Machine Learning and Instrument Autonomy: Allowing Spacecraft To Do More, With Less

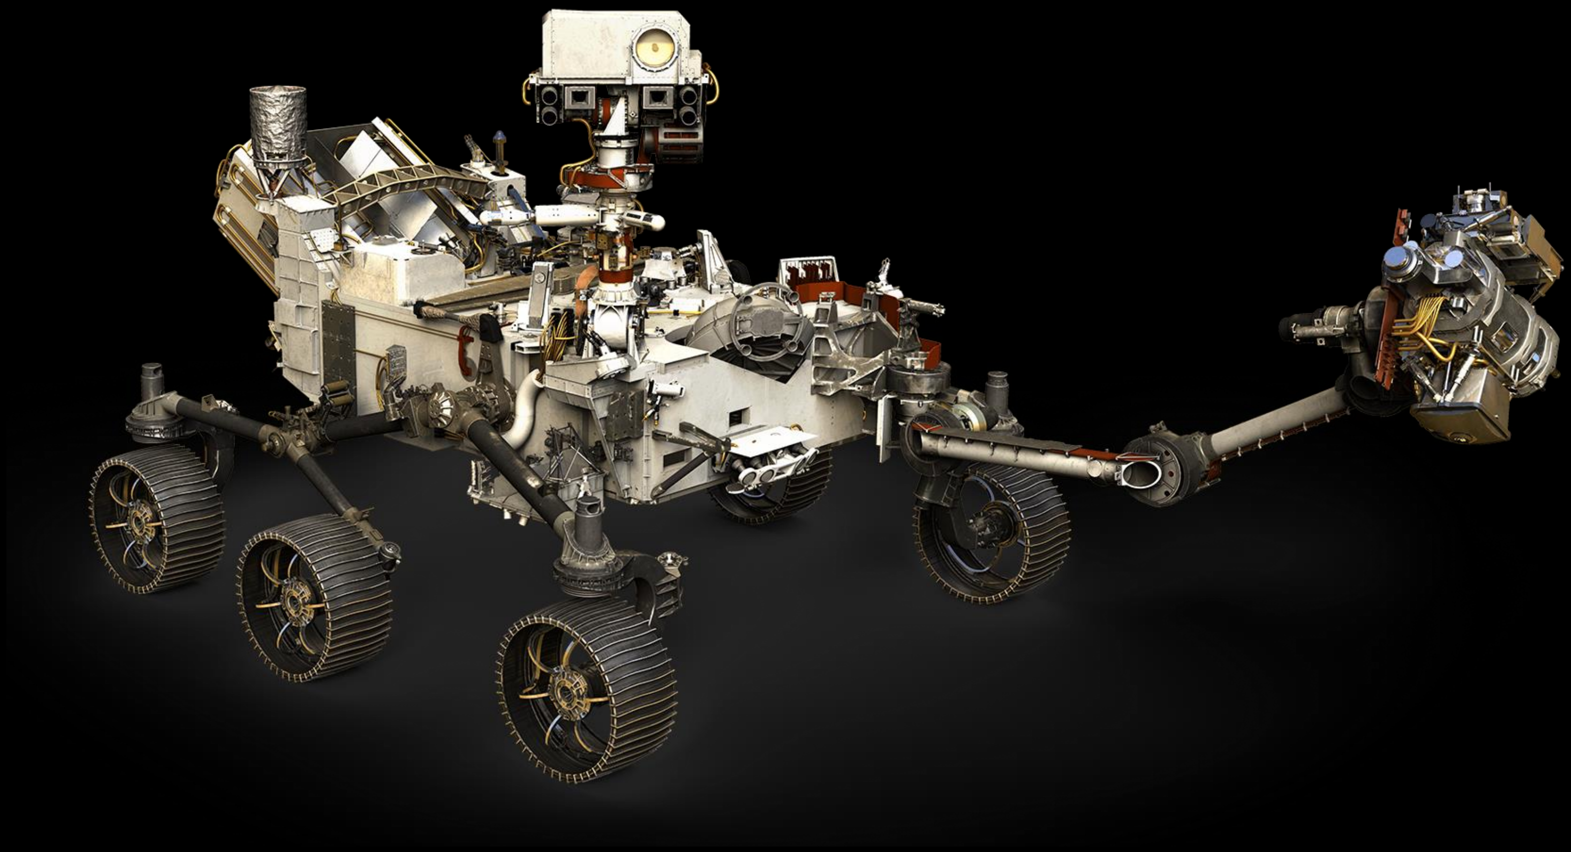

Jack Lightholder

This research was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Jet Propulsion Laboratory
California Institute of Technology

When to Use ML / Autonomy?

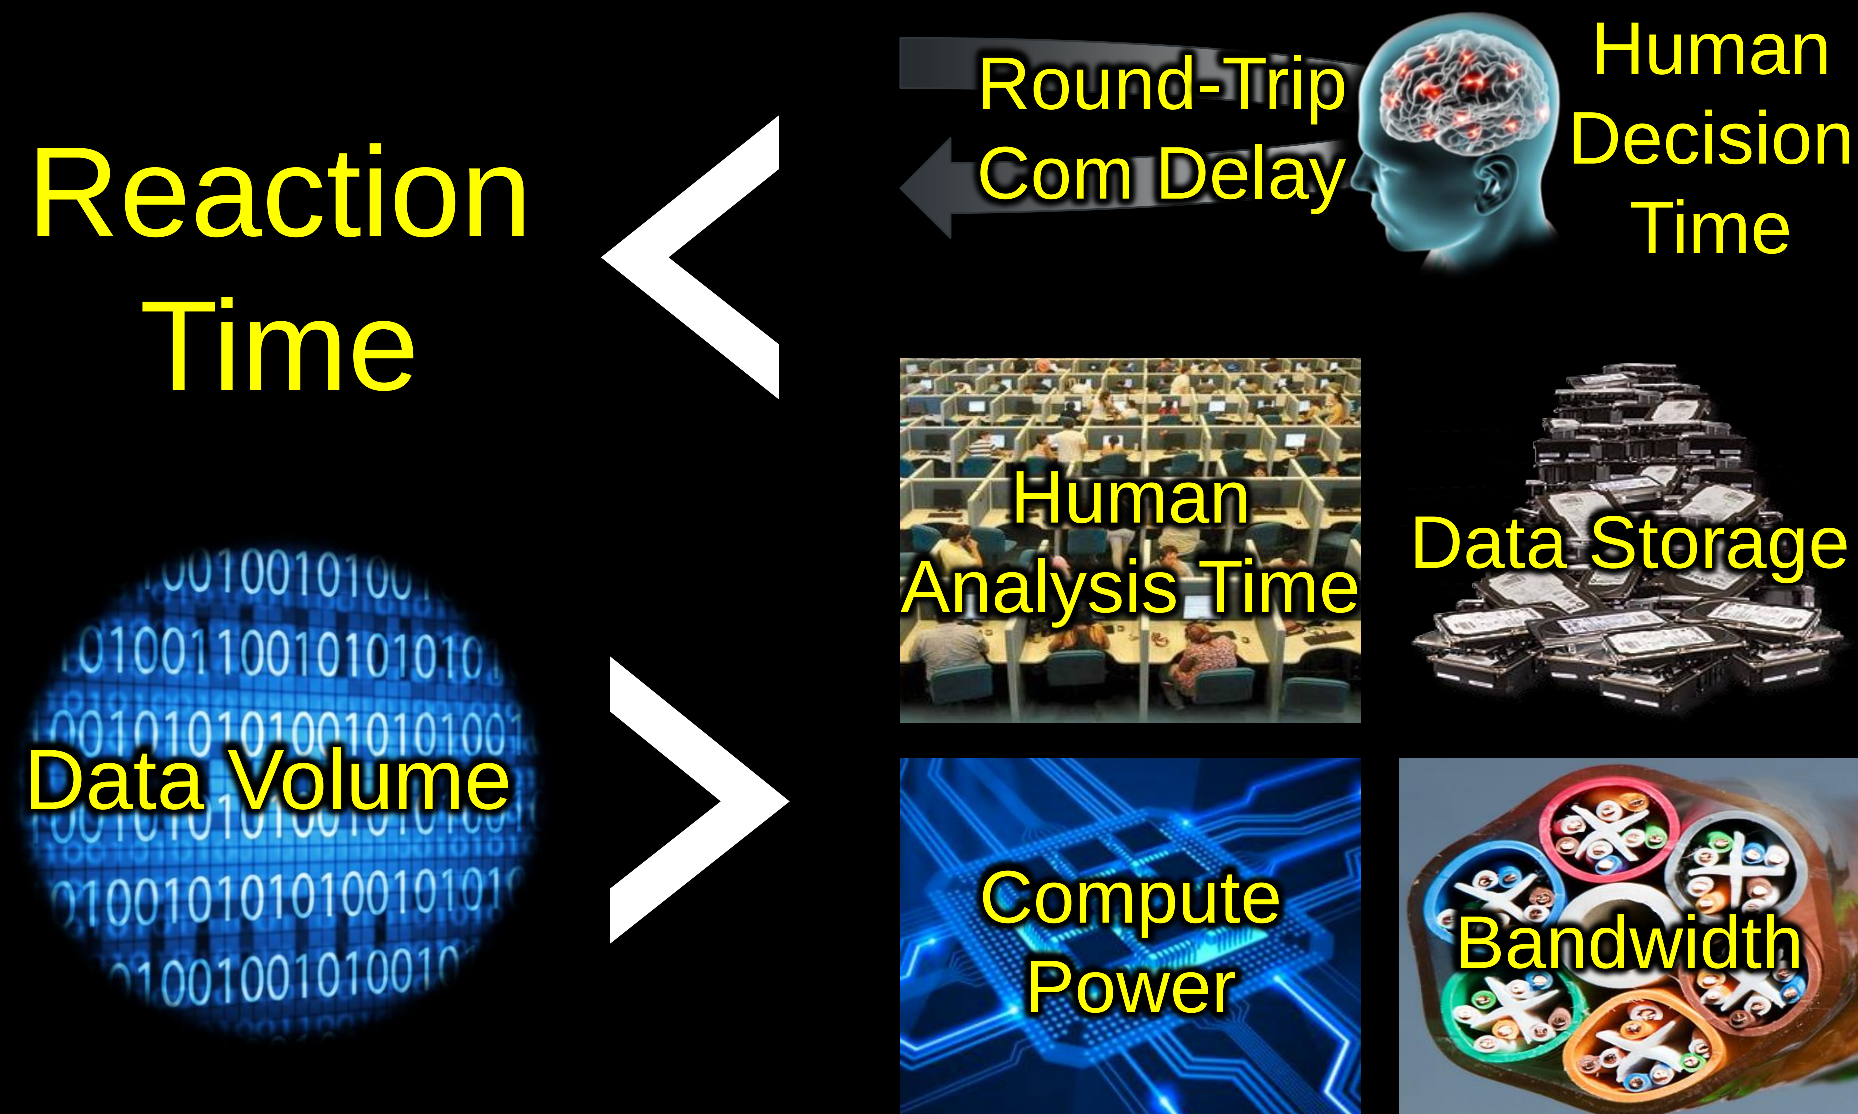

Courtesy Dr. Lukas Mandrake, JPL

ML ≠ Expert Replacement

- Eliminates drudgery
- Operates impossibly fast
- Focuses experts on interesting cases
- Enables larger human feats

Data Science asks: “Would you like to have the same output with $^{1}/_{6}$ the experts or x6 the output with your current experts?”

Courtesy Dr. Lukas Mandrake, JPL


Low Earth Orbit Examples

Detecting Features of Interest

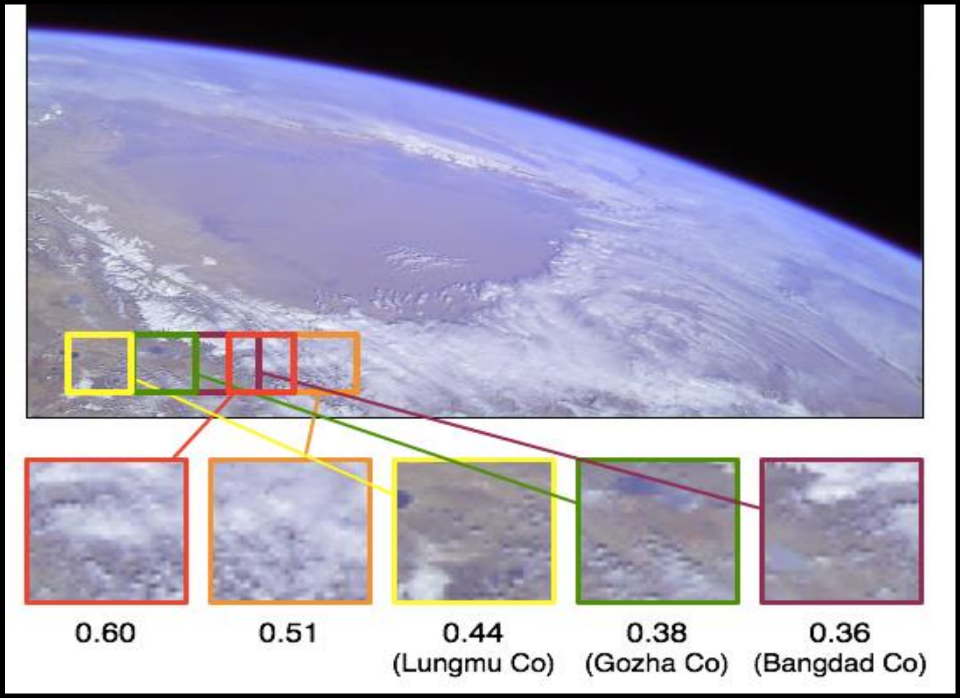

Visual Salience: Identified areas of the image that differ from surrounding areas.

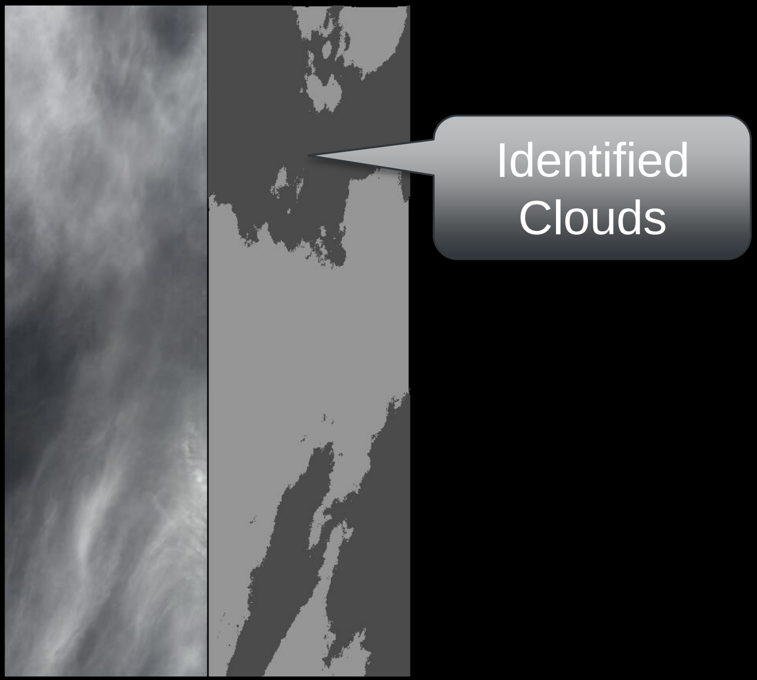

Preliminary Cloud Classification results from EO-1

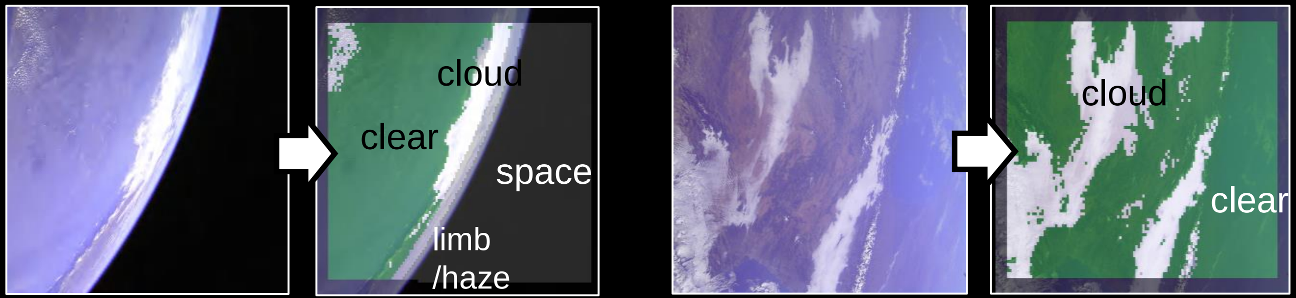

TextureCam: Pixel classification for cloud screening, downlink prioritization

[Chien et al. 2016 JAIS; Thompson et al., i-SAIRAS 2012; Wagstaff et al., *GRL* 2013;
Bekker et al., *Astrobiology* 2014; Altinok et al. JFR 2015, Thompson et al. TGARS 2011]

Current Model

On-board Processing
Observe, Thumbnails, transmit

Science Investigation
Manual Inquiry

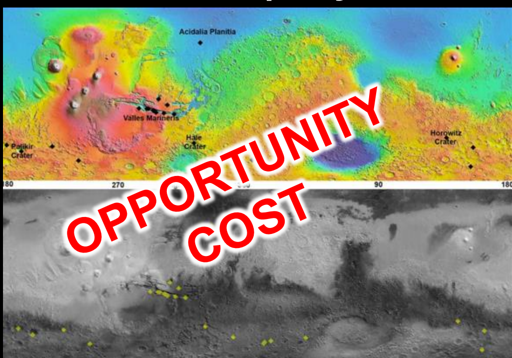

Ops Decision Support
Host of Scientists, Manual Selection

Martian Orbit
Unmapped /
Changing Surface

Courtesy Dr. Lukas Mandrake, JPL

Data-Driven Assistance

Science Support
Data Mining

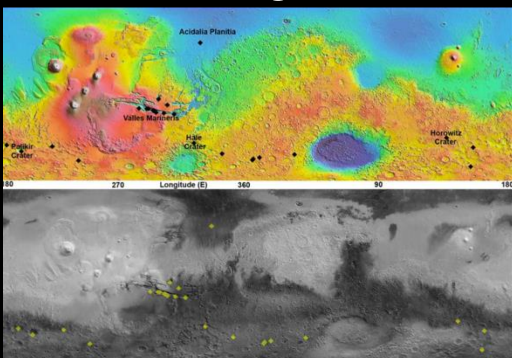

Martian Orbit
Unmapped /
changing surface

Ops Decision Support
Focus of Attention Tools

On-board Science
Detect Transients,
Summarize Content

Courtesy Dr. Lukas Mandrake, JPL

Summarization Technology

Scene-Wide Labels

Scene Feature	Present
Dunes	Yes
Barchan Dunes	No
Small Craters	Yes
Large Craters	Yes
Fresh Impacts	No
RSL	No

Terrain Classification

Terrain Type	Image %
Flat Plain	50%
Crater Slope	25%
Dune Field	10%
Ridges	15%

Landmark Identification

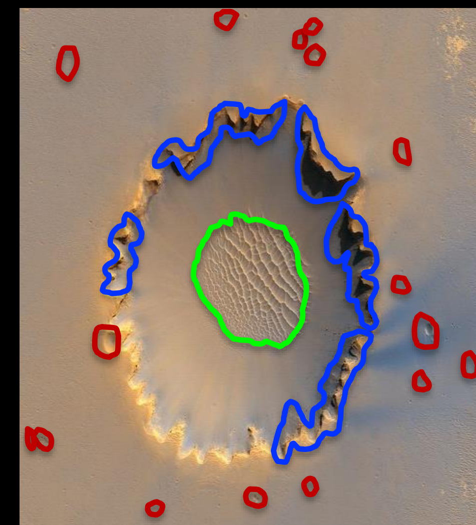

Landmark Type	Number
Small Craters	16
Ridges	4
Dunes	1

Courtesy Dr. Lukas Mandrake, JPL

Landmark Recognition

Drs. Kiri Wagstaff
Gary Doran
Ravi Kiran
Lukas Mandrake
Norbert Schorghofer
Alice Stanboli

Techniques

- **Salience Estimation**
  - Created by Genetic Algorithm
  - Finds optimal blend of leading techniques
- **Landmark Classification**
  - Naïve Bayes
  - Support Vector Machines
  - Neural Network (deep learning)

Successfully ported to:

- PDS / Planetary Image Atlas
- IPEX: Atmel 400 MHz

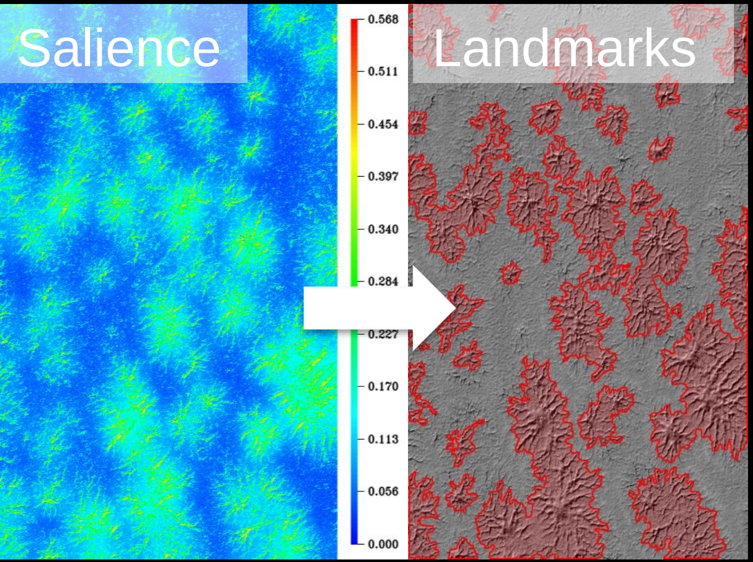

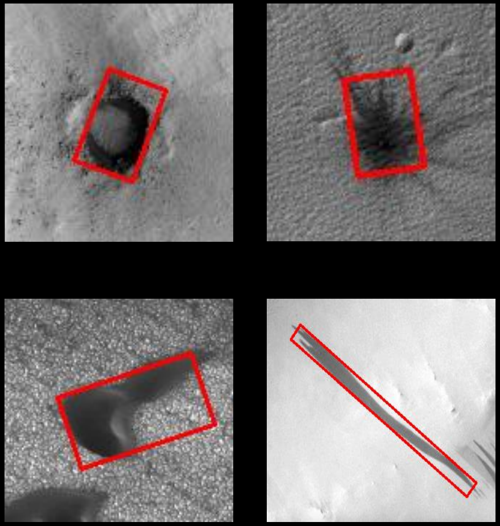

Landmark Classification

Landmark Type	Number
Small Craters	16
Ridges	4
Dunes	1

Summarization

Scene Labeling

Drs. Alphan Altinok
Brian Bue
Alice Stanboli
Kiri Wagstaff

<u>"Scalable Scene Analysis" System</u>

- Convolutional Neural Network
- Implemented on PDS Atlas
- Currently trained for Cassini & MSL Images

craters transients rings surface horizon clouds plume
sky view starfield body types multiple objects phases
artifact eclipse haze over exposure noise ripple camera distance

19 categories – 53 labels

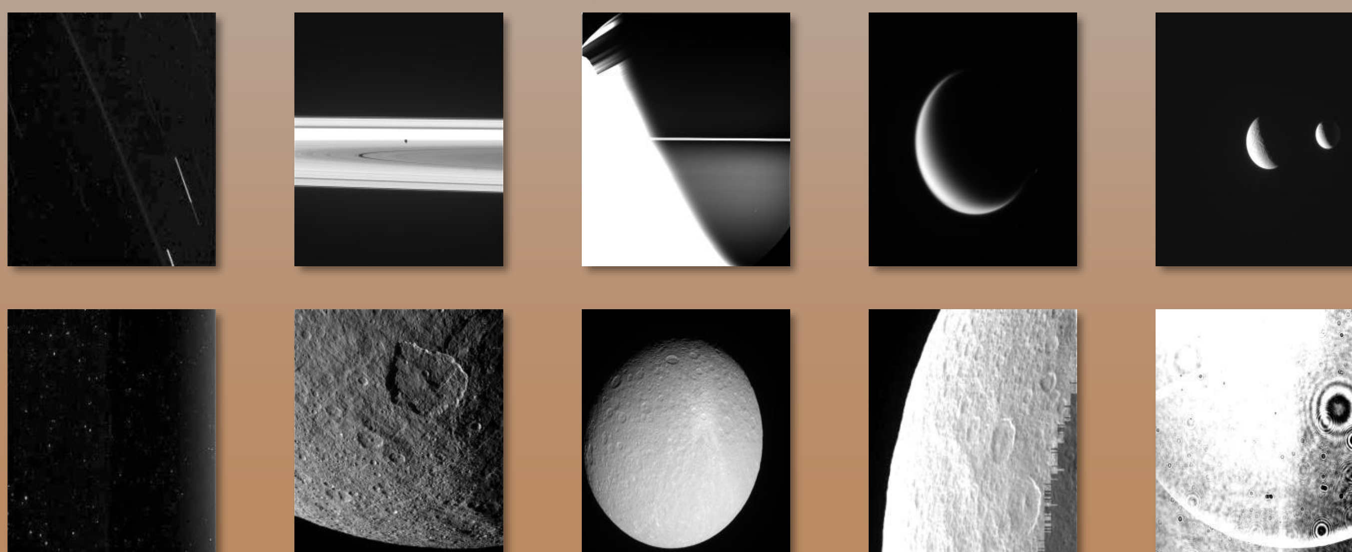

Terrain Classification

TextureCam System

- Random Forest based pixel classifier
- Extremely fast & parallelizable

Successfully ported to:
- MSL VSTB Flight Testbed (RAD750) = ~100 HiRise images/day
- EO-1: Mongoose V (M5) processor
- IPEX: Atmel 400 MHz

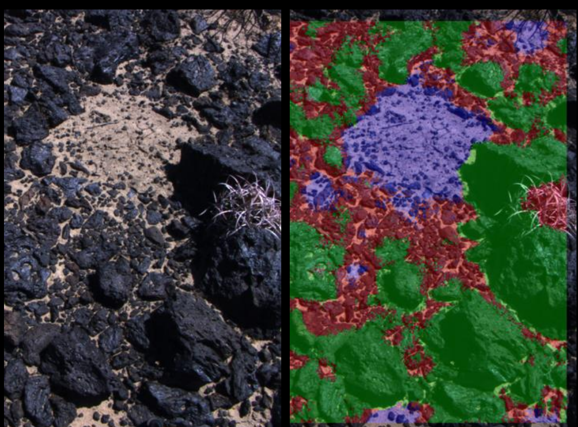

Cima Lava Fields

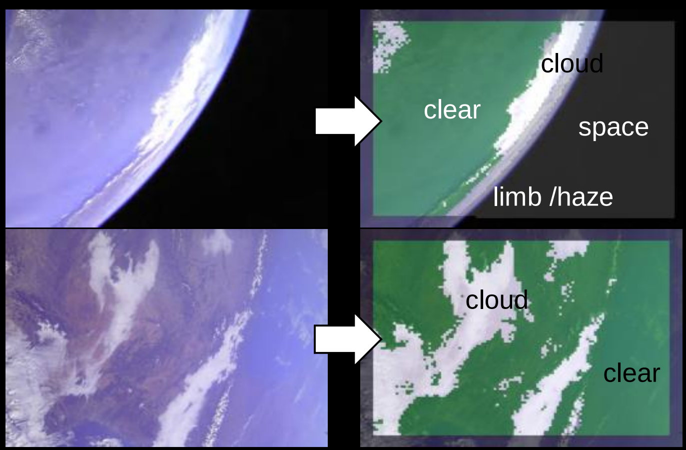

IPEX Cube-Sat Feature Identification & Cloud Mask

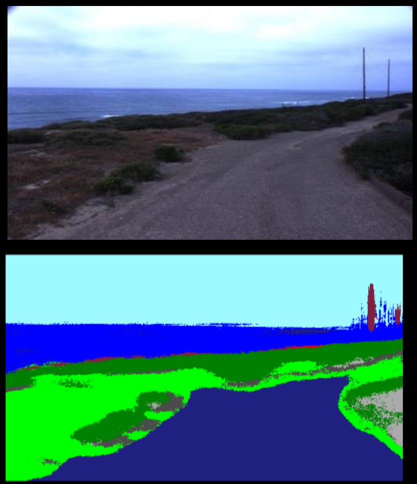

Collaboration MLIA & MV

Drs. David Thompson, Alphan Altinok, Brian Bue, Gary Doran, Kiri Wagstaff

A Landed Mission Example

AEGIS

Autonomous Exploration for Gathering Increased Science

- Target & Zap Rock
- Manually Scheduled Targets

- Round Trip Delays
- Trouble hitting 1st time
- Targeted science not possible right after drive

- Autonomy selects interesting targets
- Refines targeting automatically

- ~30-100% additional ChemCam science targets on drive sols

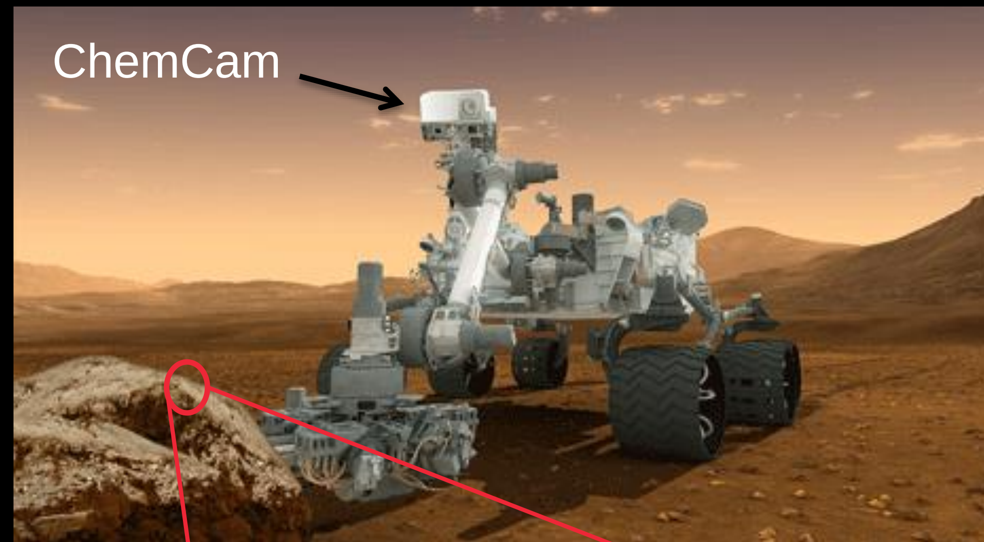

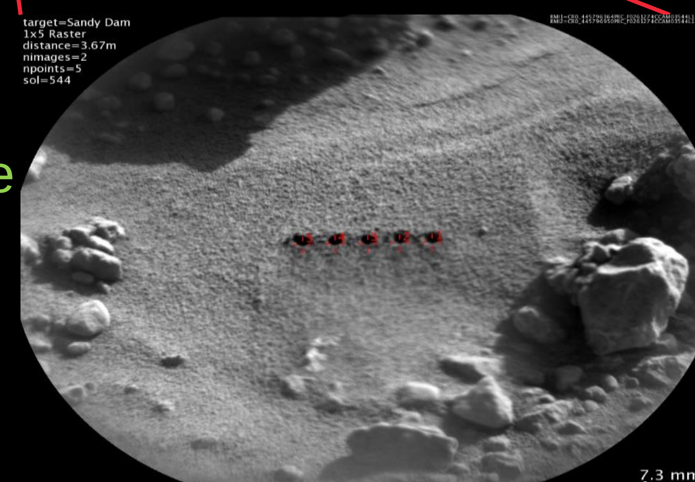

Drs. Tara Estlin, Dan Gaines, Gary Doran, Raymond Francis, et al.

Mars Exploration Rover (2009)

Mars Science Laboratory (2012)

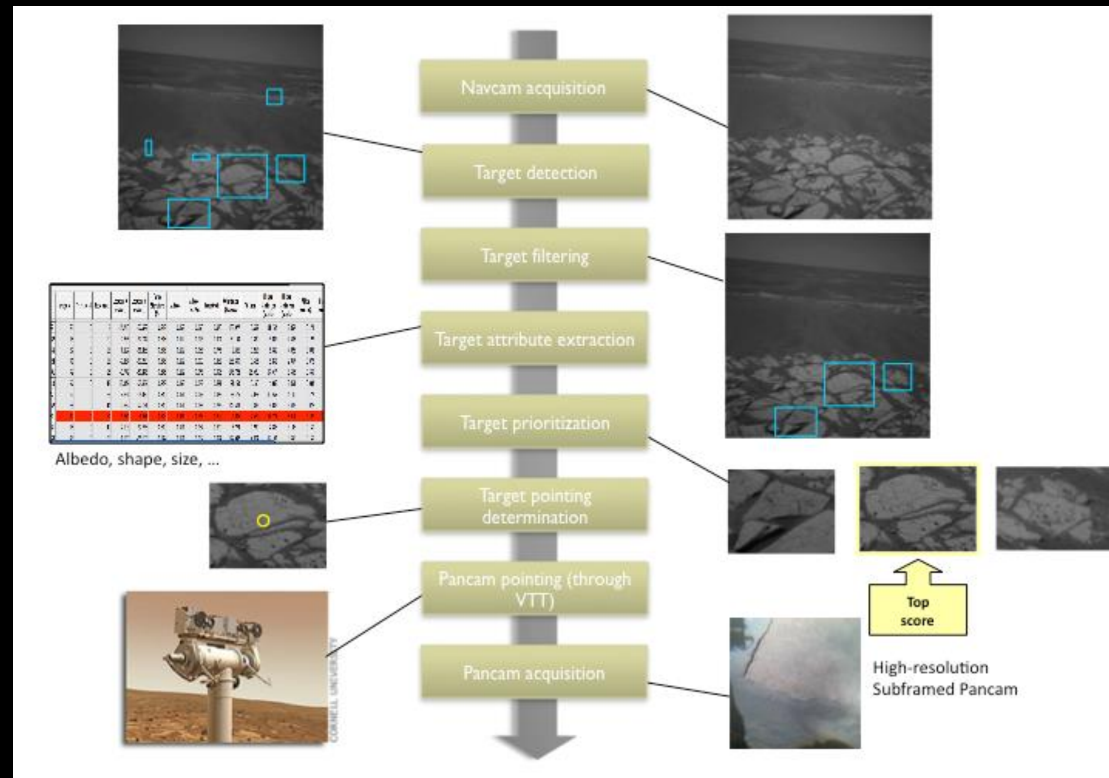

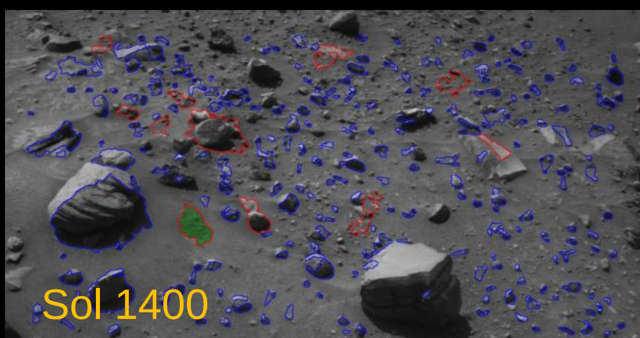

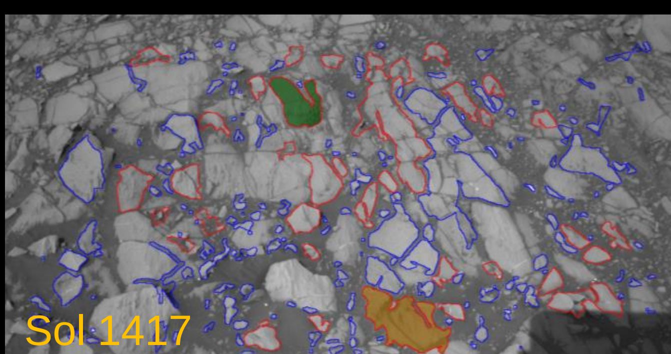

- Provides **intelligent targeting and data acquisition** by:
  - analyzing images of the rover scene
  - identifying high-priority science targets (e.g., rocks)
  - taking follow-up imaging of these targets with no ground communication required

[Estlin et al. 2011]


A Deep Space Example

Mission Agility Through Onboard Analysis

Analyze data acquired onboard spacecraft and respond based on analysis

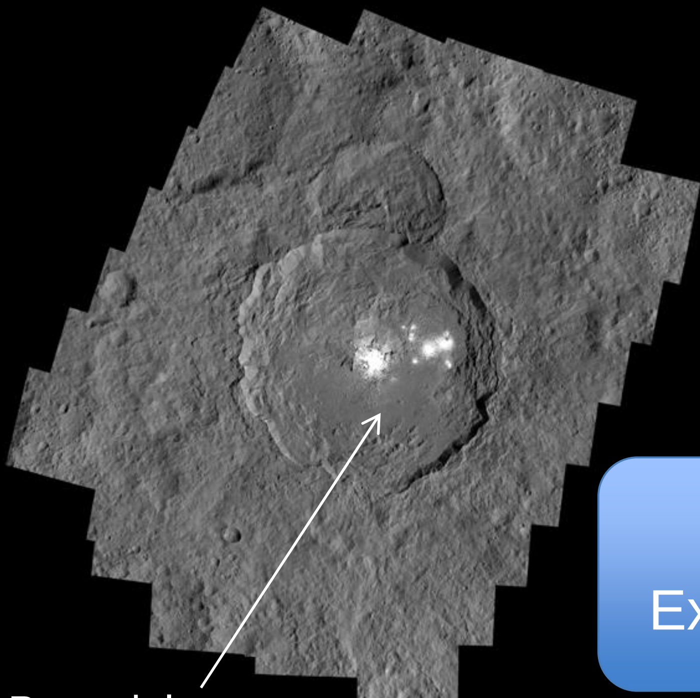

Potential volatiles on Ceres

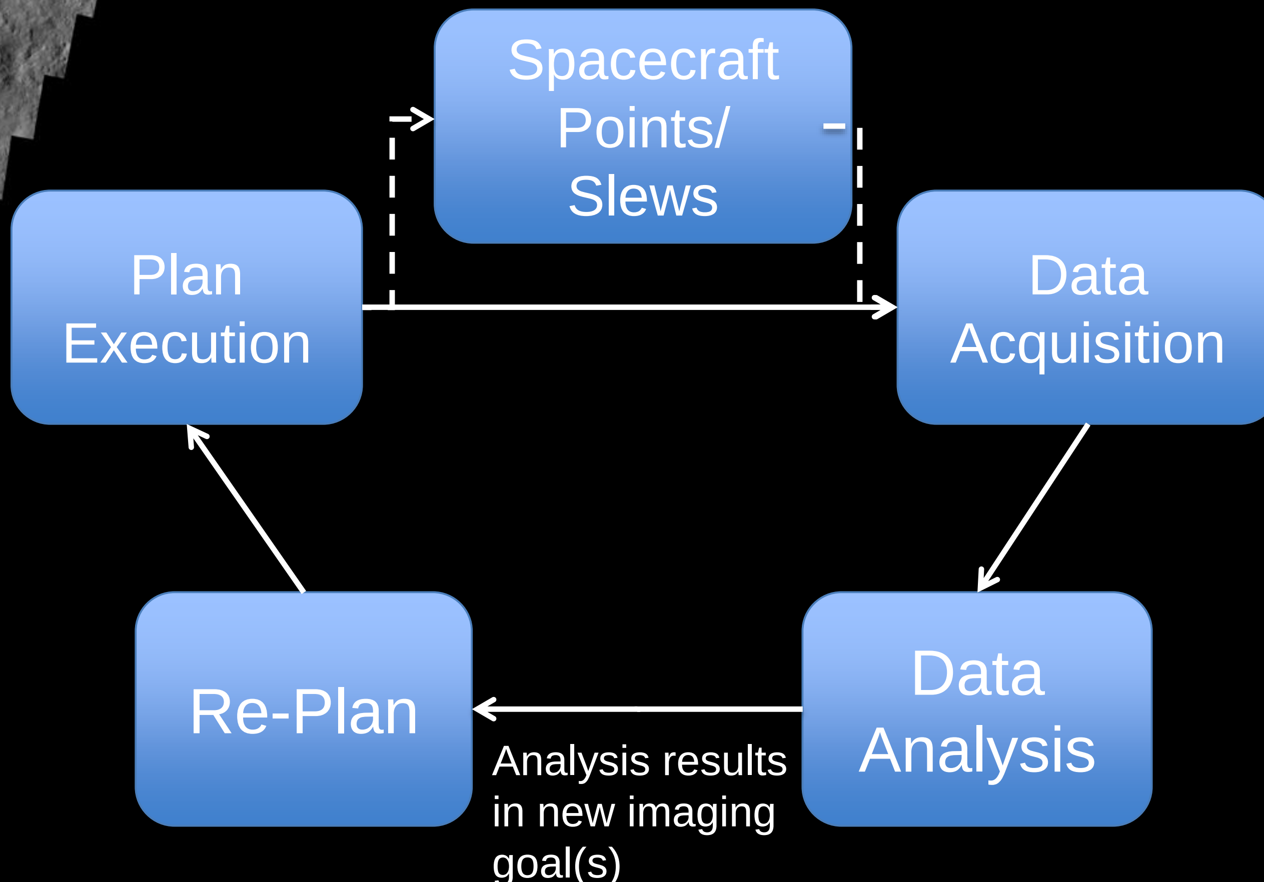

Image Credit: NASA/JPL/Dawn Mission

Near Earth Asteroid Scout

GOALS

Characterize one candidate NEA with an imager to address key Strategic Knowledge Gaps

Demonstrates low cost capability for HEOMD for NEA detection and reconnaissance

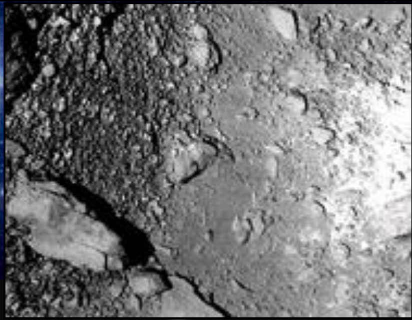

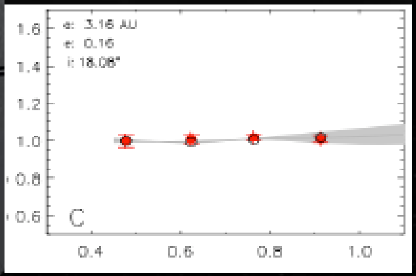

Measurements: *NEA volume, spin and orbital properties, address key physical and regolith mechanical SKGs.*

Imaging Challenges

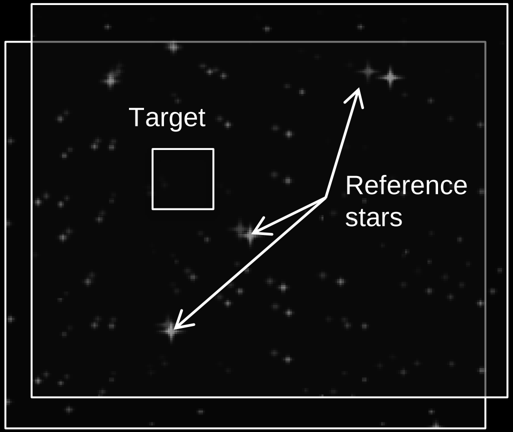

Target Detection and Approach
Ephemeris determination

Target Position Uncertainty

Spacecraft Pointing and Camera Limitations

Medium Field Imaging
Shape, spin, and local environment

Short Flyby Time (<30 minutes)

Uncertain Environment

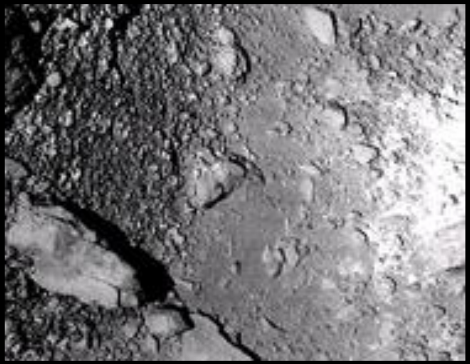

Close Proximity Imaging
Local scale morphology, terrain properties

Data Value Analysis and Sorting

Short Time at Closest Approach (<10 minutes)

Limited Downlink of 1 Kbps

Raw Data is Messy

Rosetta OSIRIS Narrow Angle Camera Detection of 2867 Steins

Raw Data is Messy

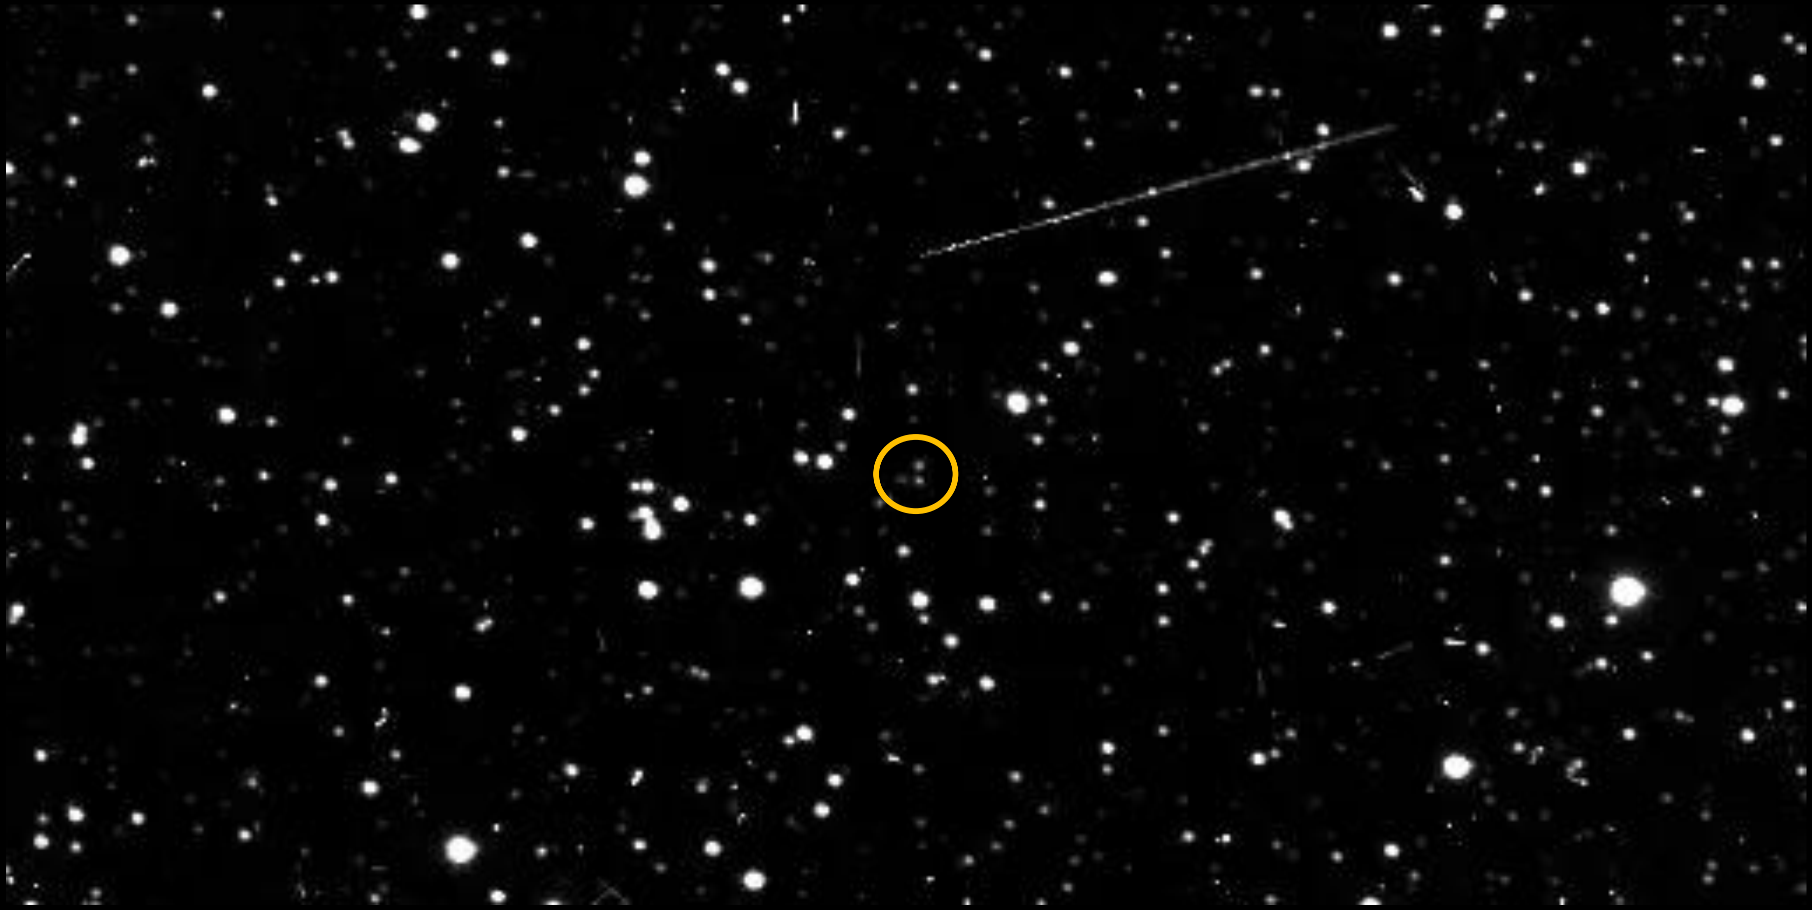

Rosetta OSIRIS Narrow Angle Camera Detection of 2867 Steins

Processed Data

Rosetta OSIRIS Narrow Angle Camera Detection of 2867 Steins

Does Your Target Look "As Expected"?

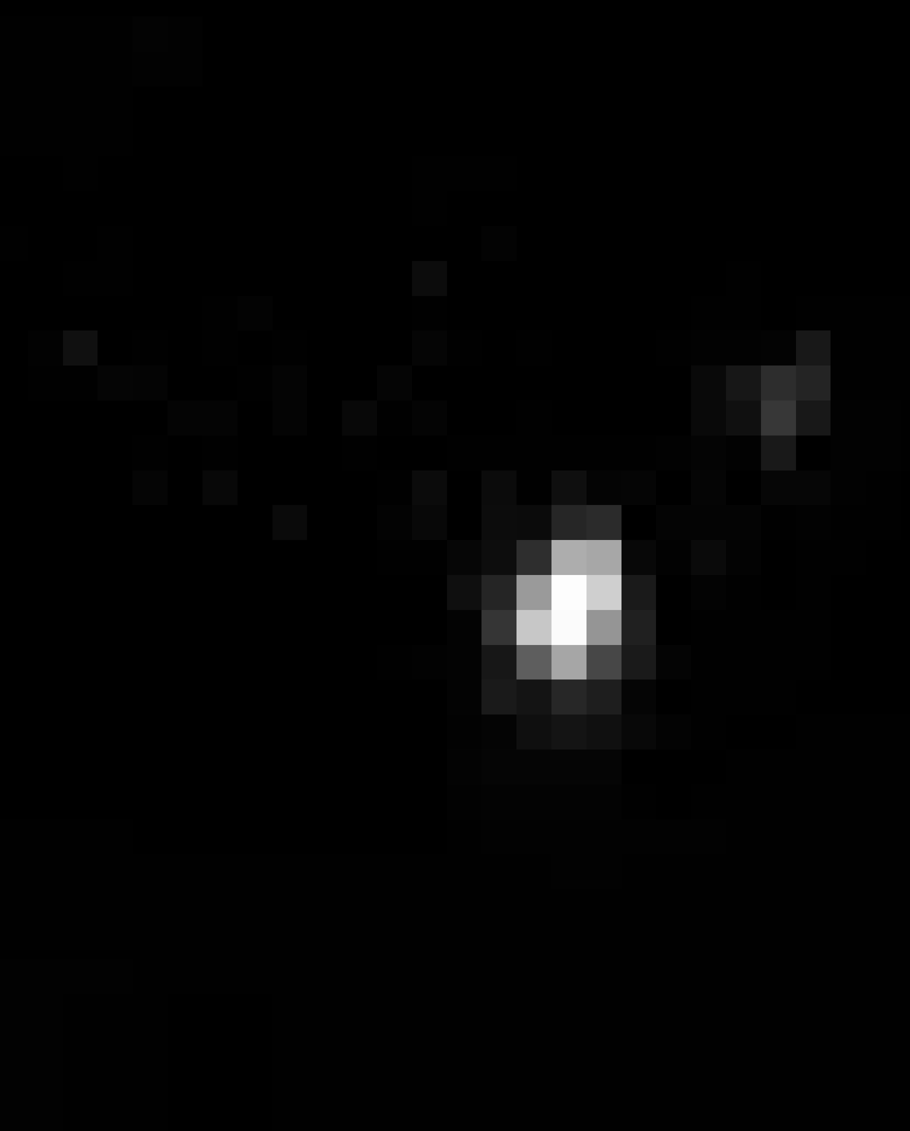

New Horizons Long Range Reconnaissance Imager Detection of Pluto / Charon

What Else Could We See?

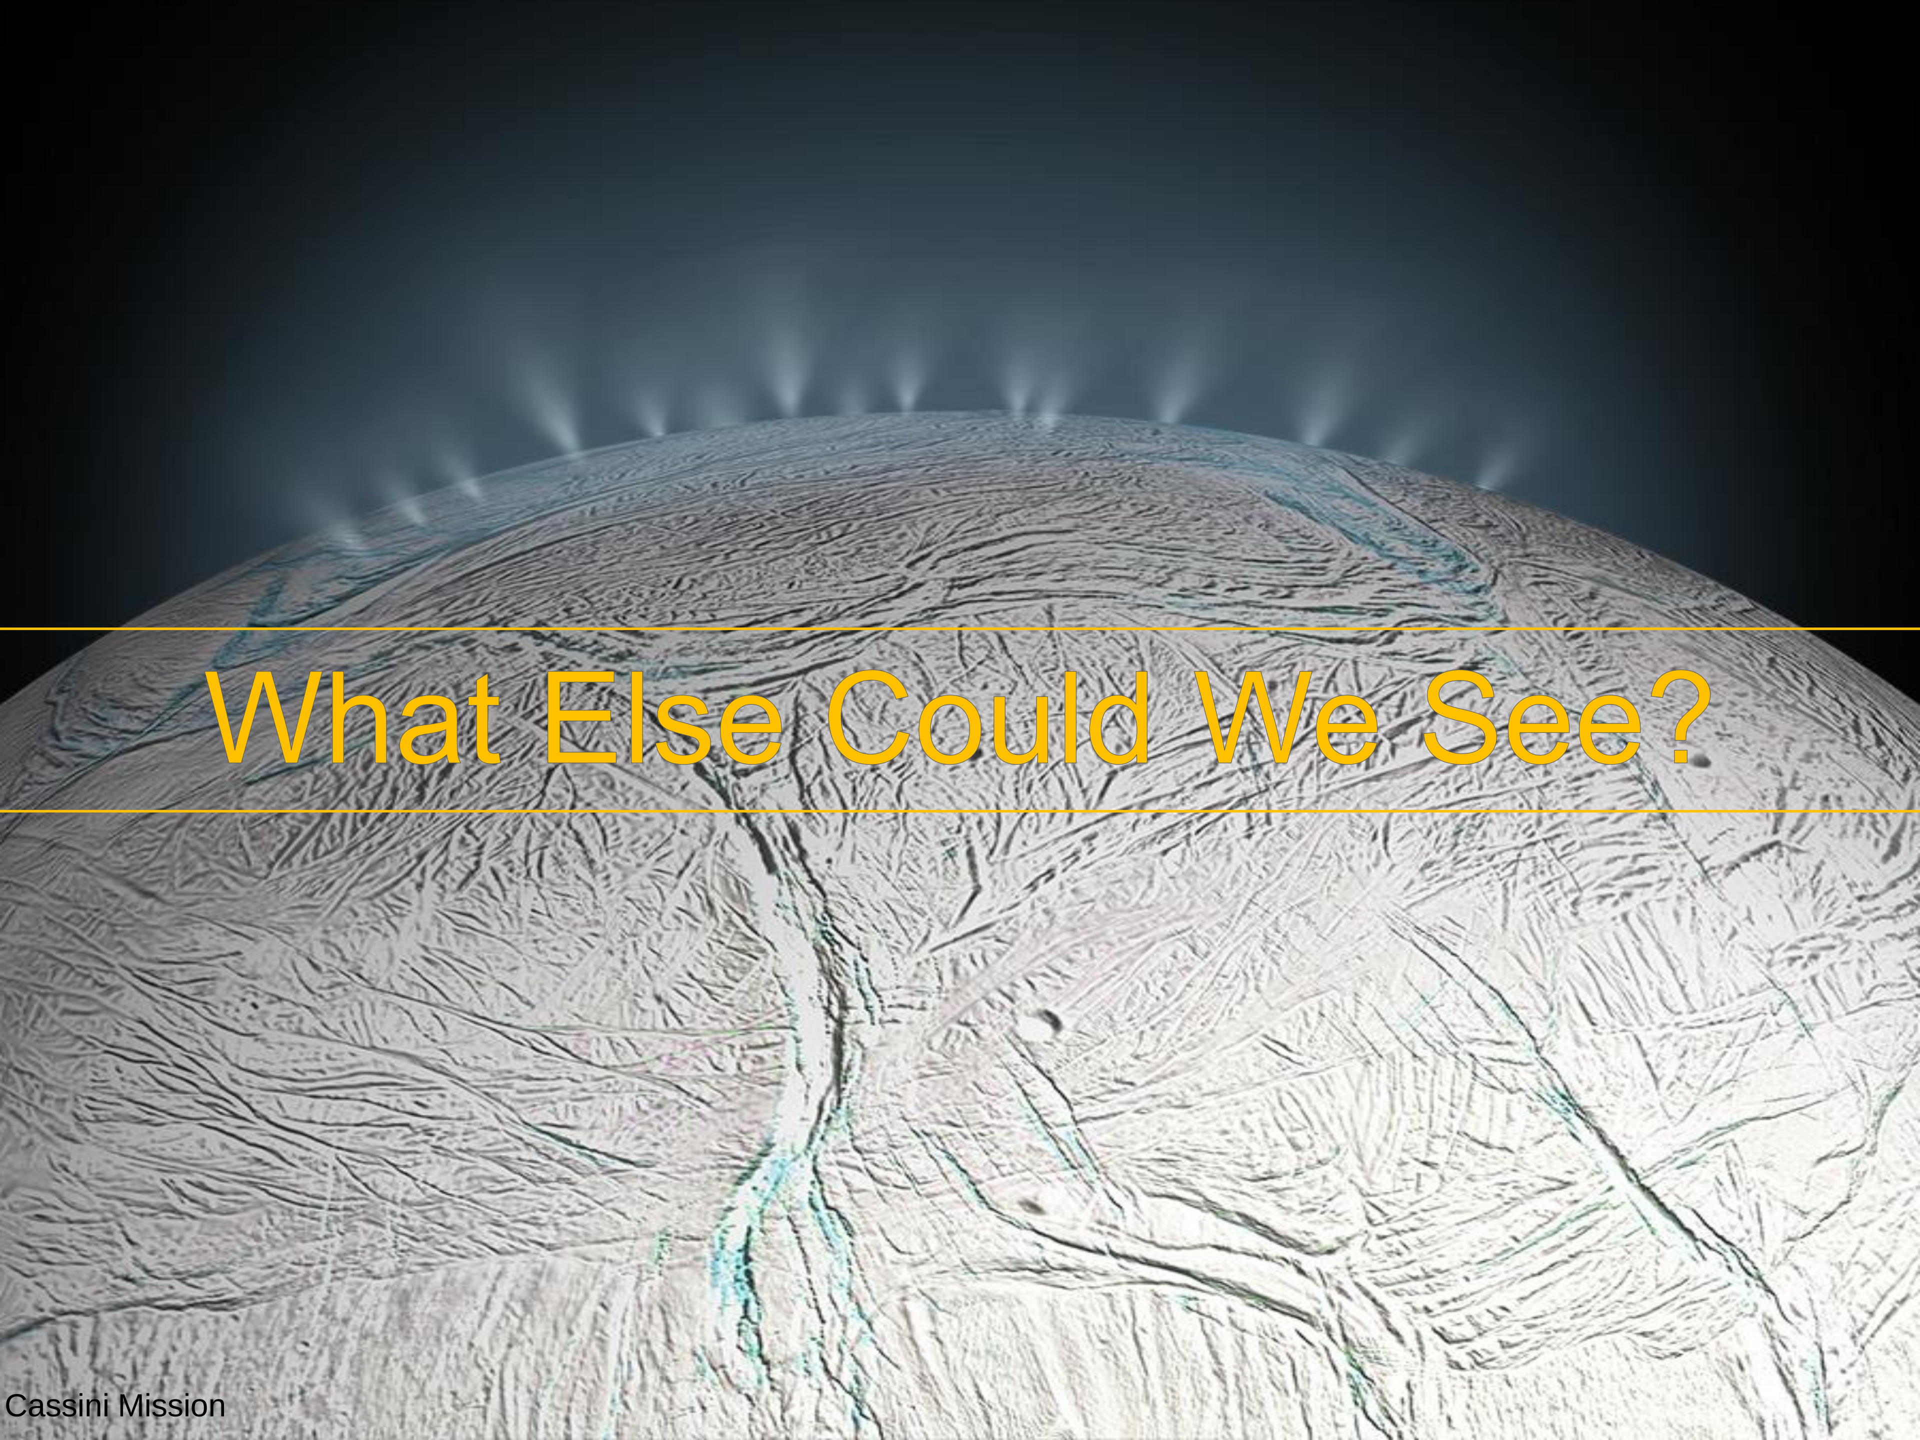

Cassini Mission

Plumes are Scientifically Exciting

Enceladus

Comet 67P

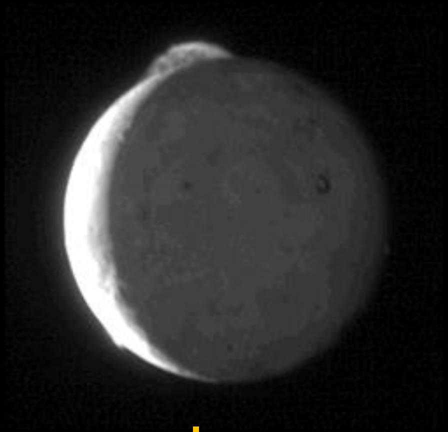

Io

Plumes gives scientists insights into the volatiles located throughout the solar system.

Unfortunately, they're not scheduled. We have to react fast.

Image Credit: NASA, Cassini Mission (Left), ESA, Rosetta Mission (Center), NASA, New Horizons Mission(Right)

Plume Detection

Hartley 2 (*EPOXI*)

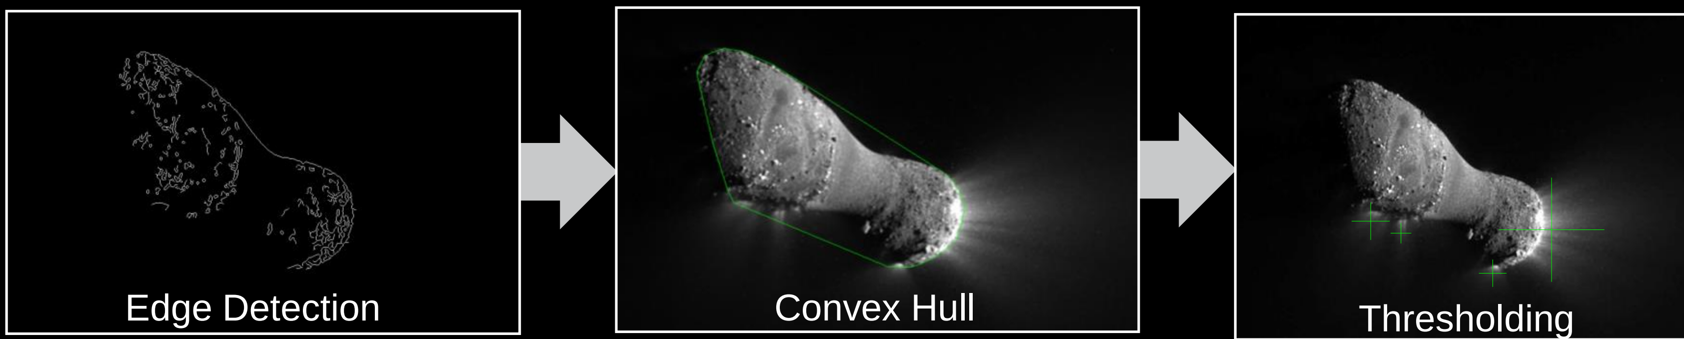

- Detects bright material beyond the limb
- Enables monitoring campaigns, target-relative data acquisition
- Detects most plumes with zero false positives

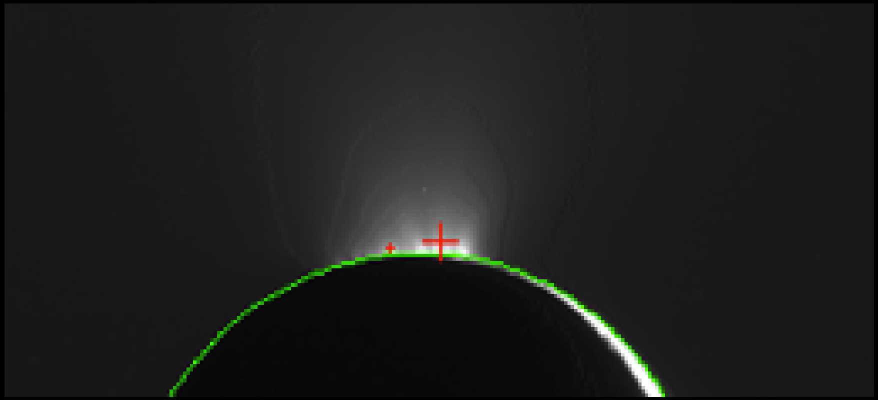

Enceladus (*Cassini*)

[Thompson et al., PSS 2012, Wagstaff et al, ApJ 2014]

Comet Tracking

Hartley 2 flyby
Original Sequence

Agile Science Planning

Target and plume detection using MRI-VIS on EPOXI. Tracking Comet Hartley 2

MOSAIC: Mars on-site shared analytics information and computing

Understand and maximize the effect of HPSC on Mars exploration

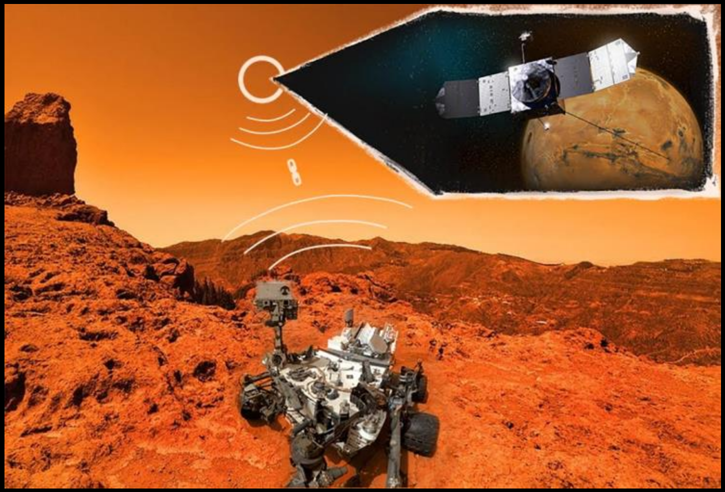

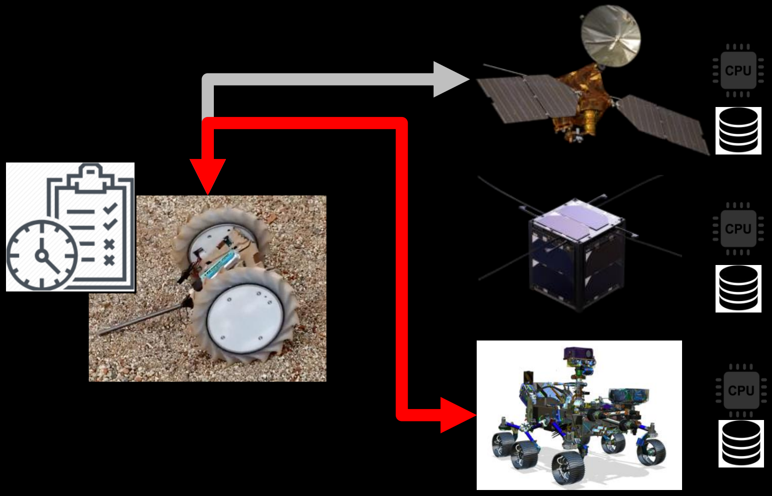

Goals:

1. Distributed computing for Mars
2. Quantify HPSC impact on missions
3. Explore trade space of HPSC designs

Research Tasks

- Resource-aware process scheduling across a network of agents
- Model-based flight computing configuration for multi-processor / multi-robot systems
- Optimize routing and storage of information across a network of agents
- Extend Delay / Disruption tolerant networking for use in distributed systems

Joshua Vander Hook (347N) hook@jpl.nasa.gov 818-354-5455

MOSAIC

1. Develop responsive, model-driven distributed computing stacks

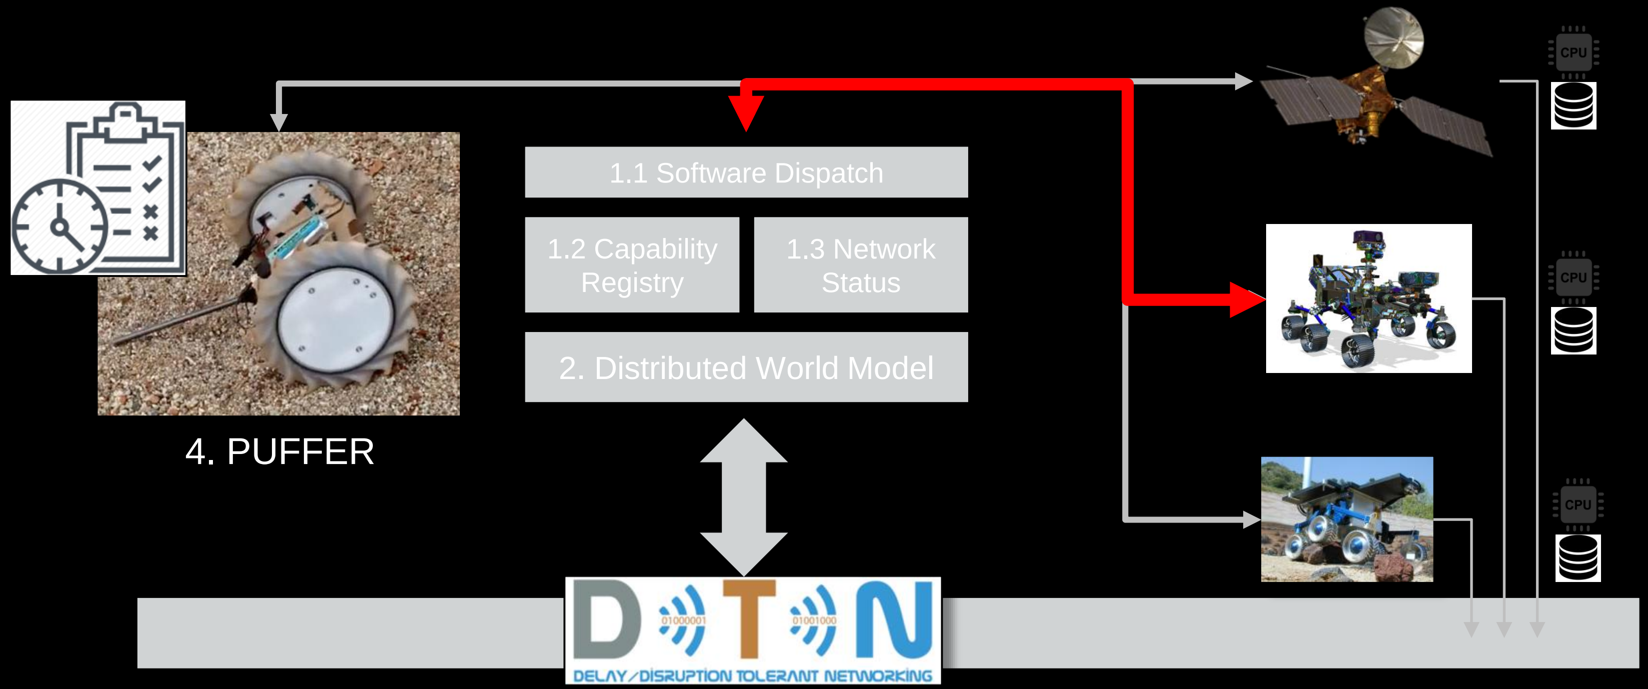

Tasks

- Benchmark existing flight software on a variety of computing hardware
- Develop analytical models to estimate runtime, data, energy requirements *as a function of HPSC config*
- Develop distributed process dispatcher (load balancing) based on above models
- Develop distributed data product consensus over DTN

Joshua Vander Hook (347N) hook@jpl.nasa.gov 818-354-5455

MOSAIC

1. Develop responsive, model-driven distributed computing stacks

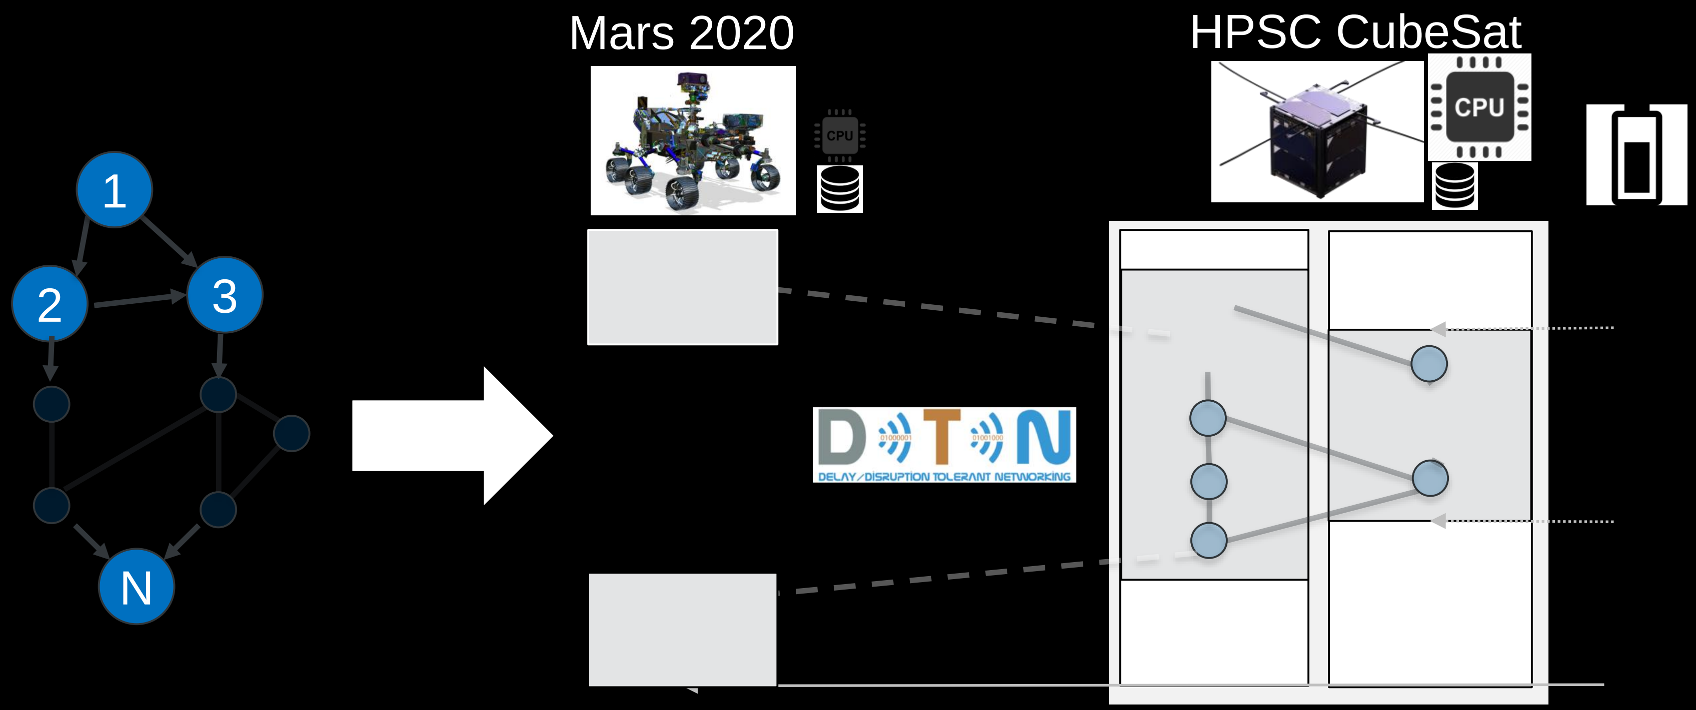

Working Example:

- Can optimally solve Mars 2020 fast-traverse FSW allocation, given HPSC + network configuration
- Output: minimum-cost allocation (time, power, etc)
- See: `"Dynamic Shared Computing Resources for Multi-Robot Mars Exploration" i-SAIRAS, 2018`

Joshua Vander Hook (347N) hook@jpl.nasa.gov 818-354-5455

MOSAIC

2. Understand impact of HPSC configurations and design on missions

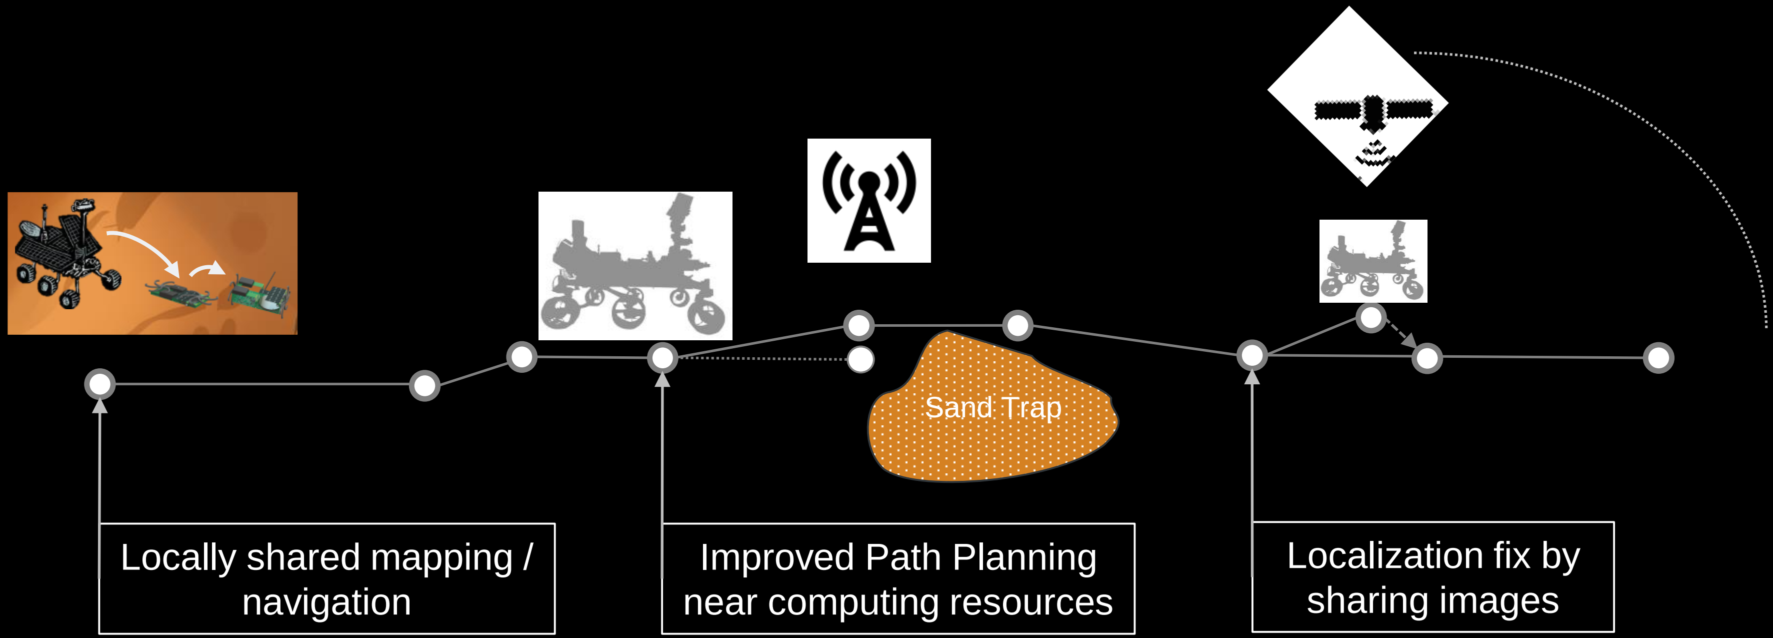

Tasks

- Given HPSC configuration, solve optimal schedule (previous) to get runtime, data, energy requirements
- Then, simulate effects on candidate missions

Joshua Vander Hook (347N) hook@jpl.nasa.gov 818-354-5455

MOSAIC

Worked example for Mars 2020 rover mission

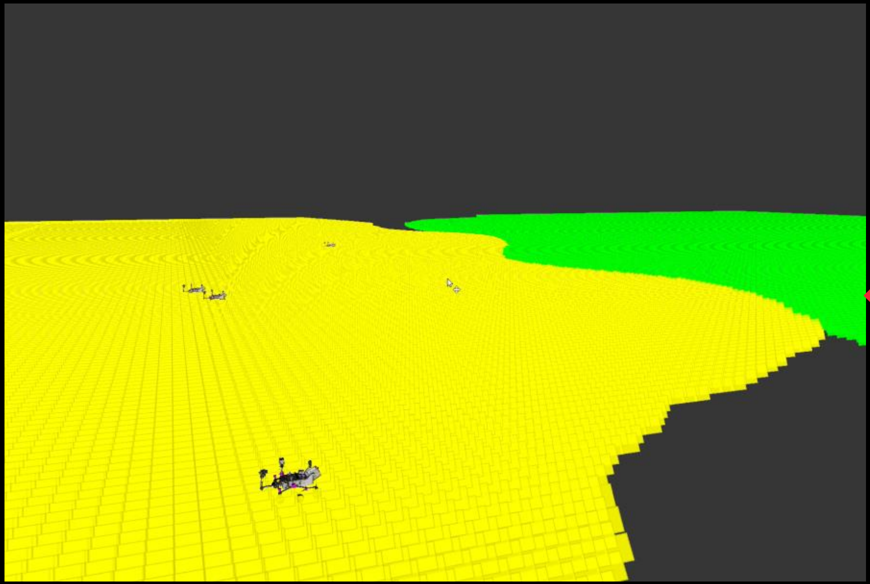

4 hardware design points, path replayed in 3D

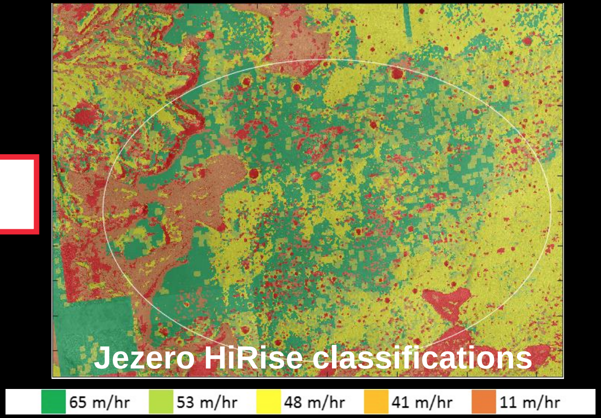

[1] Data-Driven Surface Traversability Analysis for Mars 2020 Landing Site Selection, Ono et. al.

Mars 2020 is reaches its destination 19% sooner driving through Jezero crater when it has access to three or four cores of an HPSC, either onboard, or nearby with >=1 Mbps data rate.

- Main gains are from better path optimization and better sensing
- Secondary gains from decreased sensing and planning time required

Joshua Vander Hook (347N) hook@jpl.nasa.gov 818-354-5455

MOSAIC: Mars On-Site Shared Analytics Information and Computing

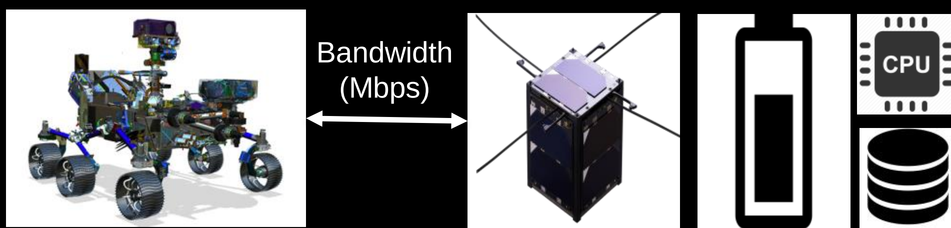

What should this hardware look like?

Given this mission

Methodology

- Given prior models, iteratively “sample” HPSC / network config to evaluate metrics
- Where possible, use “shadow cost” to determine choke point
  - (e.g., data transfer, communication bandwidth, onboard storage, or asymptotic runtime)
- Not in isolation! Consider FSW algorithms, models of environment, etc.

Joshua Vander Hook (347N) hook@jpl.nasa.gov 818-354-5455

MOSAIC

3. Explore trade space of networked multi-processor configurations

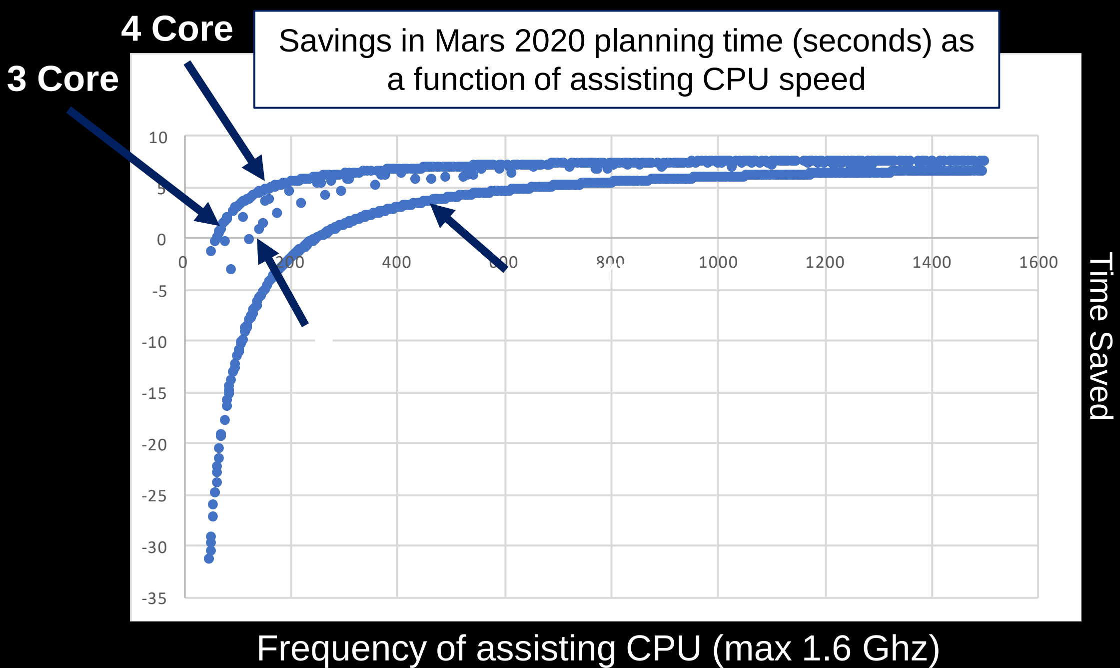

Optimal Software Process Assignment: {not possible, Rover, Assisting CPU}

Data Rate Mbps	Avg time sec	Image	DEM	Analyze Terrain	Plan Path	Drive
0.01	27					
0.1	27					
0.15	29.3					
0.3	25.6					
0.4	29.7					
0.9	28.2					
1	27.3					
100	15.3					

From Mars 2020 analysis:

- **Main gains are parallelization (3,4 core is mostly level), even at low (8%) availability**
- **Bottleneck is data rate, solution space is “level” w.r.t. compute**

Joshua Vander Hook (347N) hook@jpl.nasa.gov 818-354-5455

Science Data Prioritization

NAVCAM
~Gbytes/Sol

O(1000) images/Sol)

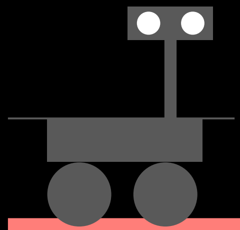

Storage
Gbytes ~ Tbytes

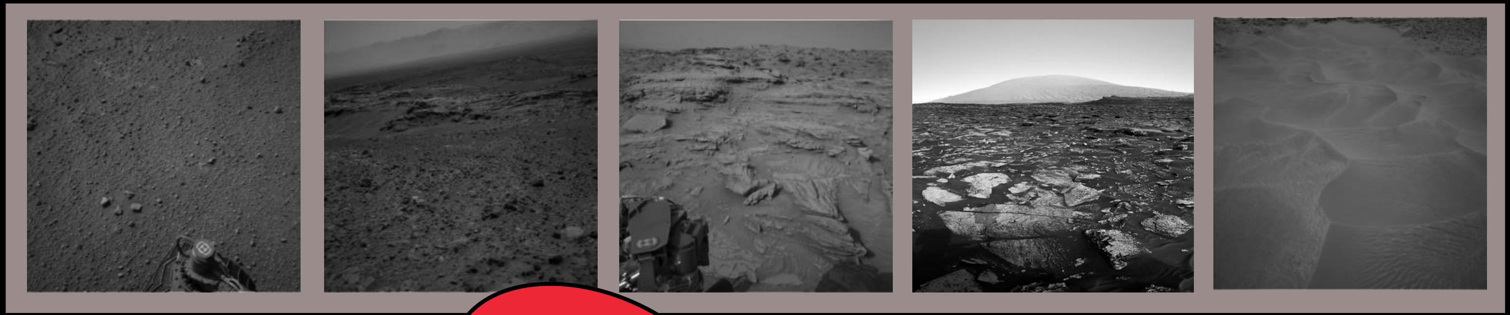

Downlink

Bonus Science

Engineering Data

0 - 100MBits/Sol (Fetch rover)

Masahiro Ono (347F) marahiro.ono@jpl.nasa.gov 818-354-0930

Energy Optimal AutoNav Preliminary Result

Energy optimal vs time optimal

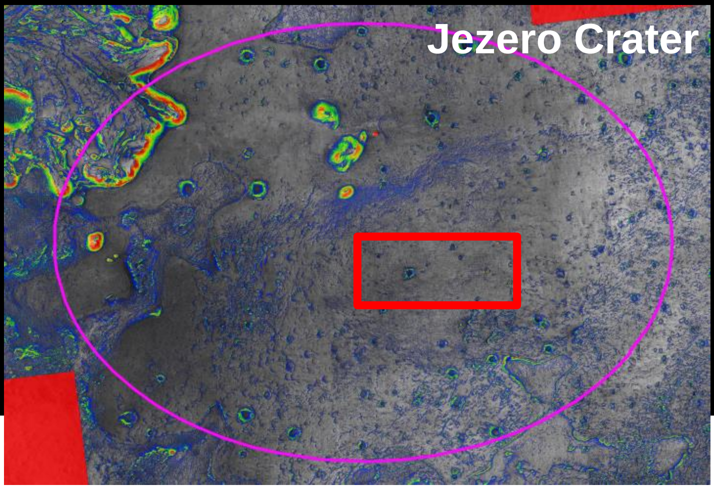

- Used Jezero Crater's DEM and terrain data
- Simulation based on Fetch Rover design
  - Performed in collaboration with Austin Nicholas
  - Used Fetch Rover's solar panel area, battery size, min charge level, nominal driving energy
  - Used MSL's slip curve
  - Used MER/InSight's dust accumulation model; assumed 100th Sol
  - Sun elevation > 10 deg
  - M2020 driving speed

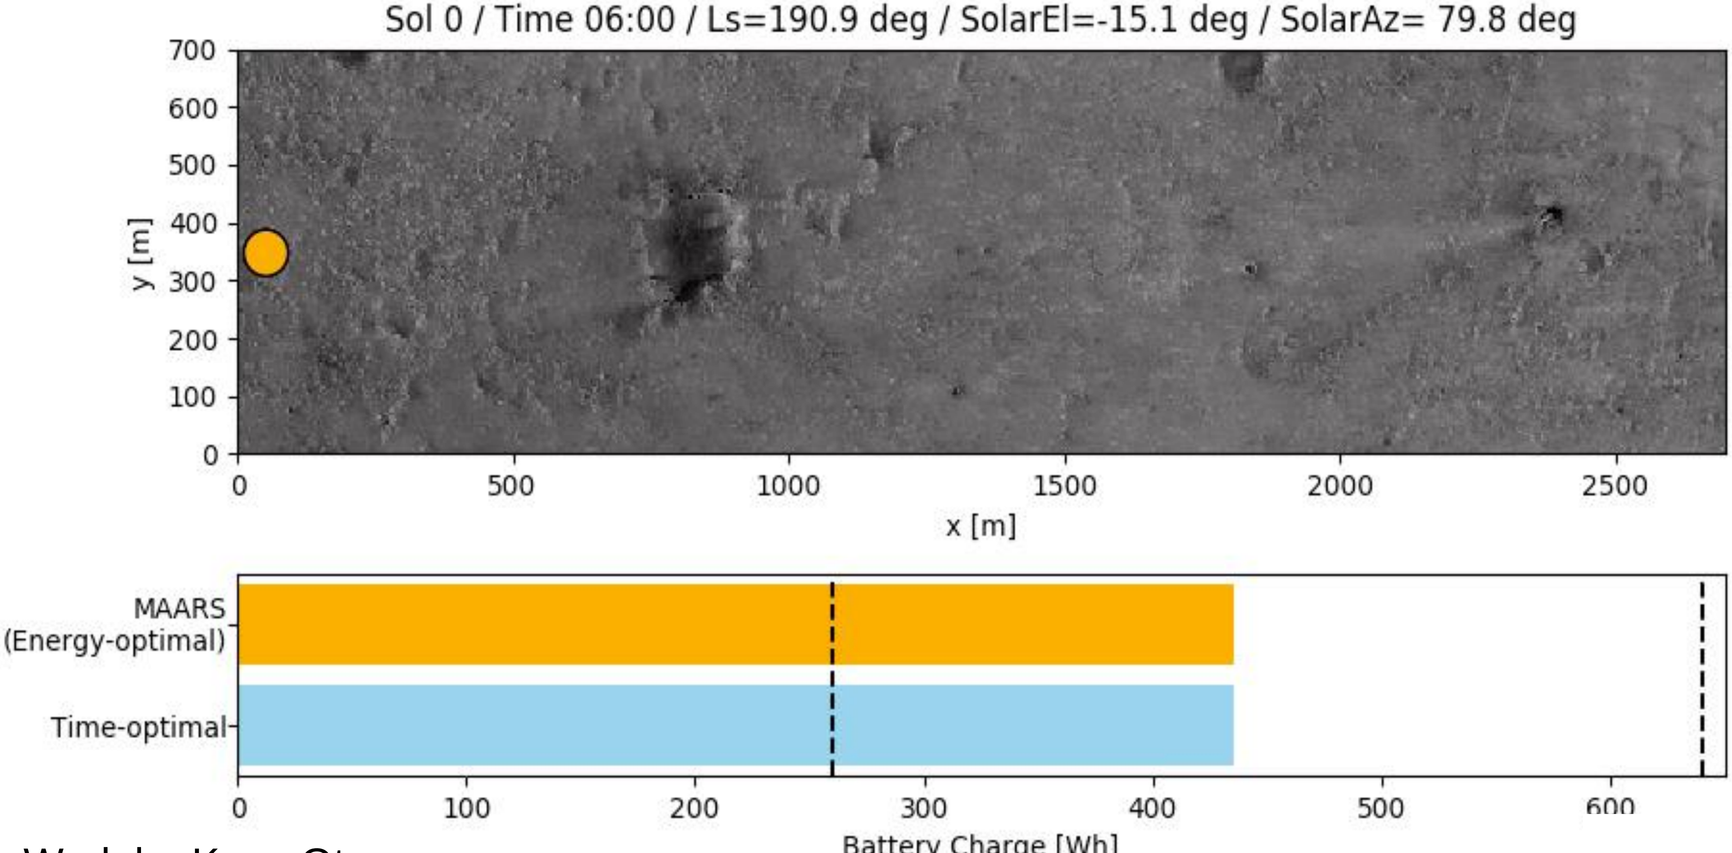

Work by Kyon Otsu

New Way of Commanding AutoNav

M2020: Command by waypoints

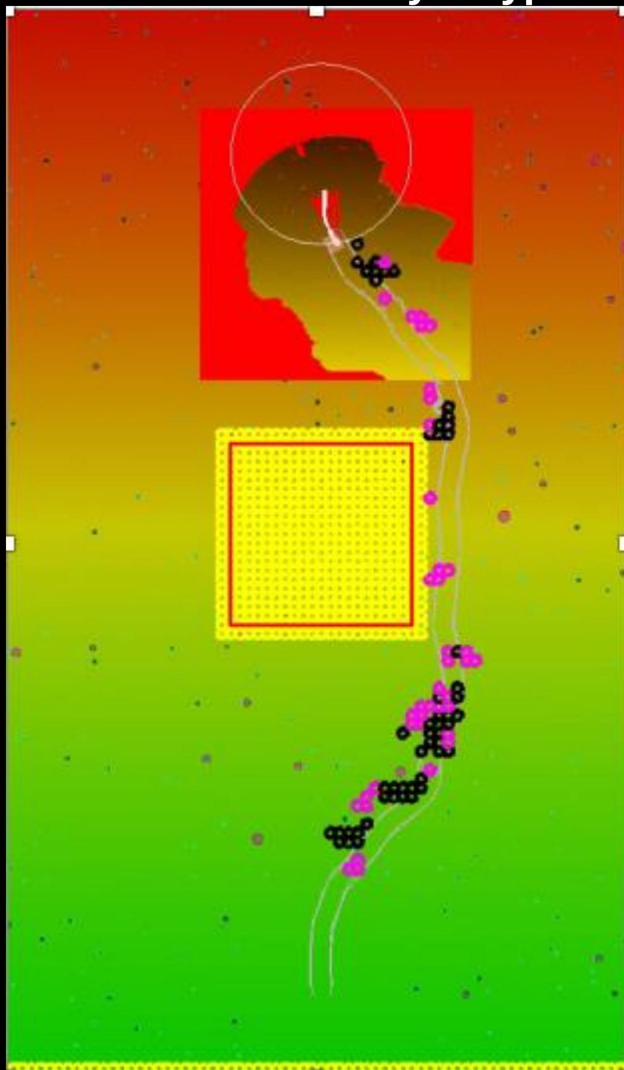

- Uplink waypoint and KOZs only
- Plan min-time path to waypoint

MAARS: Command by costmap

34	33	33	32	31	30	30	31	32	33
35	34	34	33	32	31	31	32	33	34
36	35	35	34	33	32	32	33	34	35
37	36	36	35	34	33	33	34	35	36
38	37	37	36	35	34	34	35	36	37
39	38	38	37	36	35	35	36	37	38
40	39	39	38	37	36	36	37	38	39
		40	39	38	37	37	38	39	40
				39	38	38	39	40	41
				40	39	39	40	41	42
47	46	44	43	41	40	40	41	42	43
47	46	44	43	41	41	41	42	43	44
48	47	45	44	42	42	42	43	44	45
47	45	44	43	43	43	43	44	45	46
48	46	45	44	44	44	44	45	46	47
49	47	46	45	45	45	45	46	47	48
50	48	47	46	45	46	46	47	48	49
51	49	48	47	47	47	47	48	49	50
52	50	49	48	[illegible]	48	48	49	50	51
53	51	50	49	49	49	49	50	51	52

- Uplink global cost-to-go map
  - Cost to the strategic goal from each cell
- Min local cost + global cost-to-go

Work by Kyon Otsu

Concurrent Path Planning & Scheduling

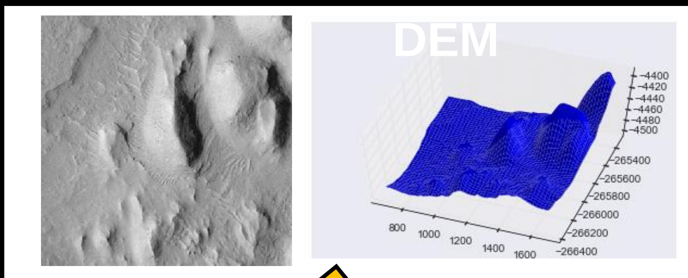

Sun simulation
(Mars24 by NASA GISS)

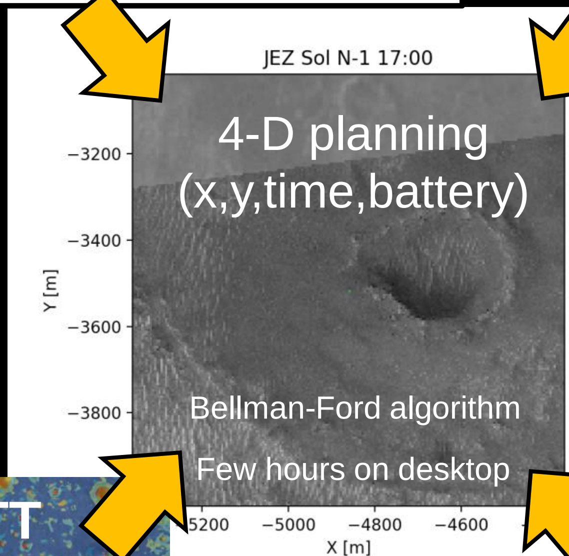

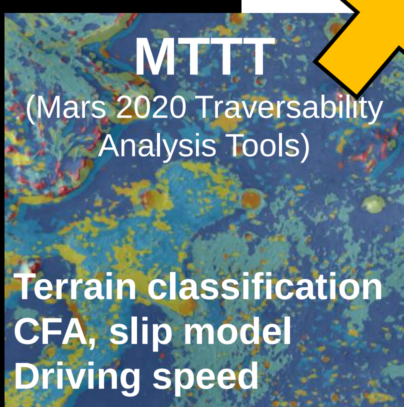

Work by Kyon Otsu

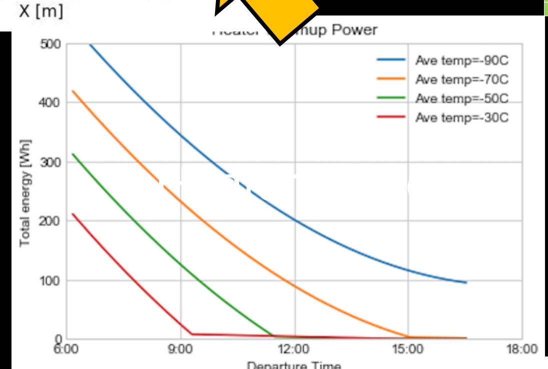

Provided by Austin Nicholas

Ground | **Mars**

Uplink

Compressed costmaps

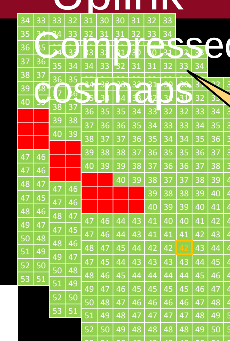

Real-time execution on HPSC

Trade off:

Data size v.s. on-board computation, performance

Preliminary Planning Results

Seasonal variation

- Used Jezero Crater's DEM and terrain data
- Simulation based on Fetch Rover design
  - Sun elevation > 10 deg
  - M2020 driving speed

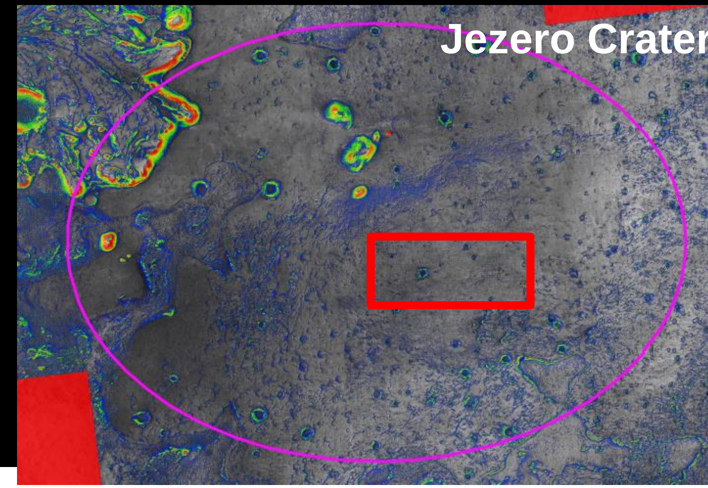

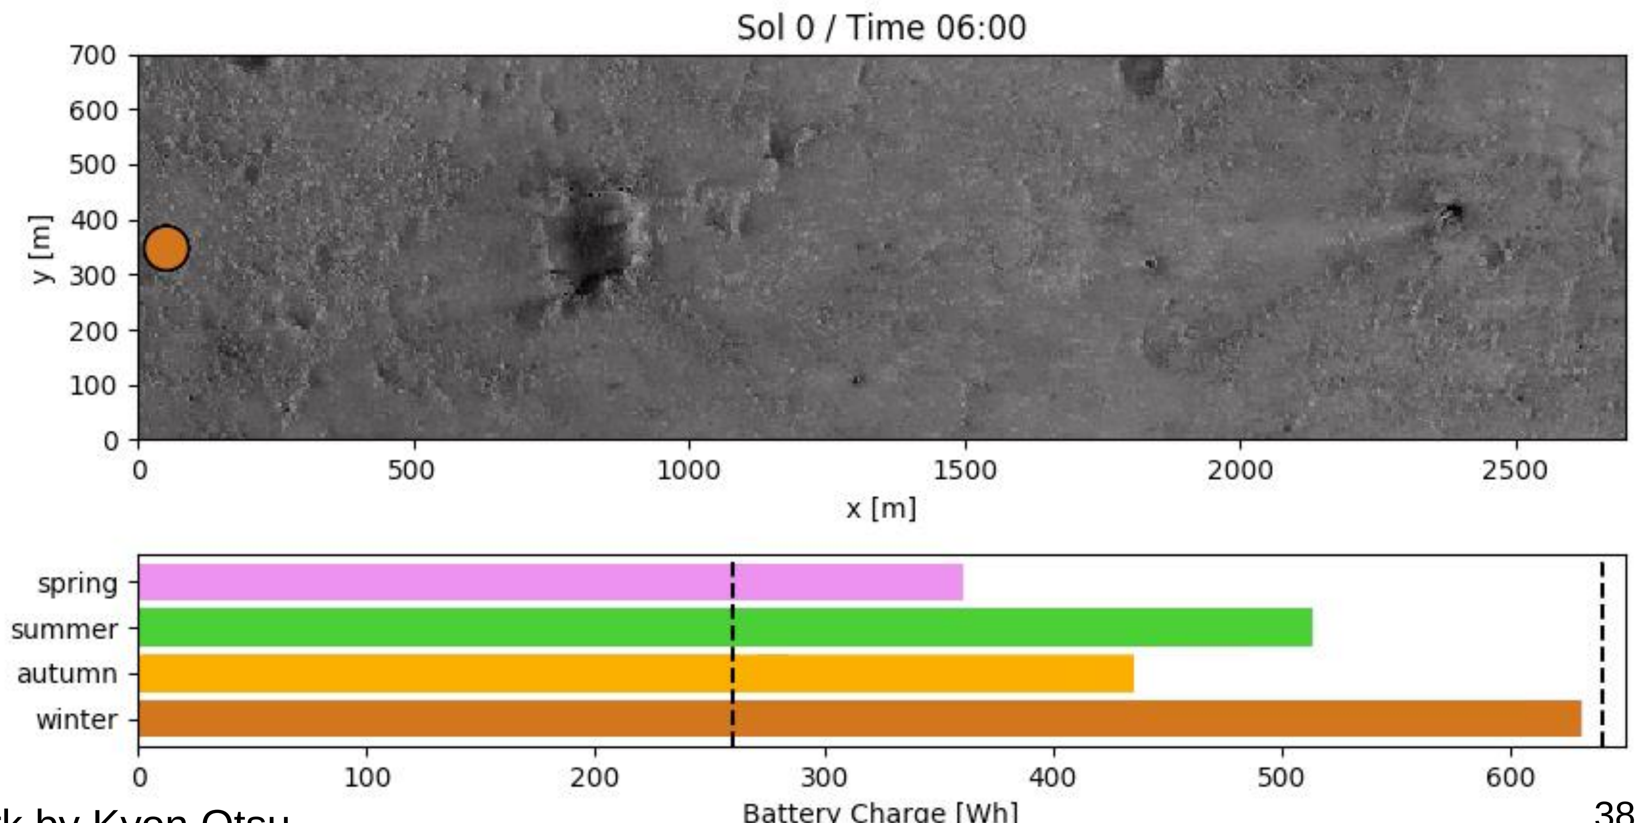

Work by Kyon Otsu

Vision-based Classification: Data Collection by Athena

Current sensor

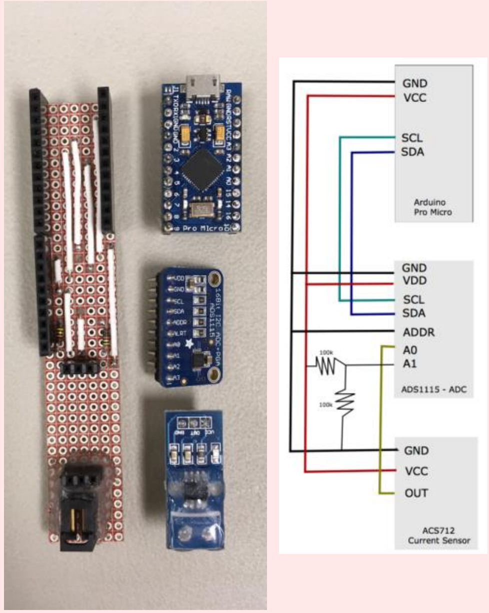

Work by Federico Zechini, Jacek Sawoniewicz, Kyon Otsu

Energy-based Terrain Classification

Clustering by HDP-HSMM algorithm

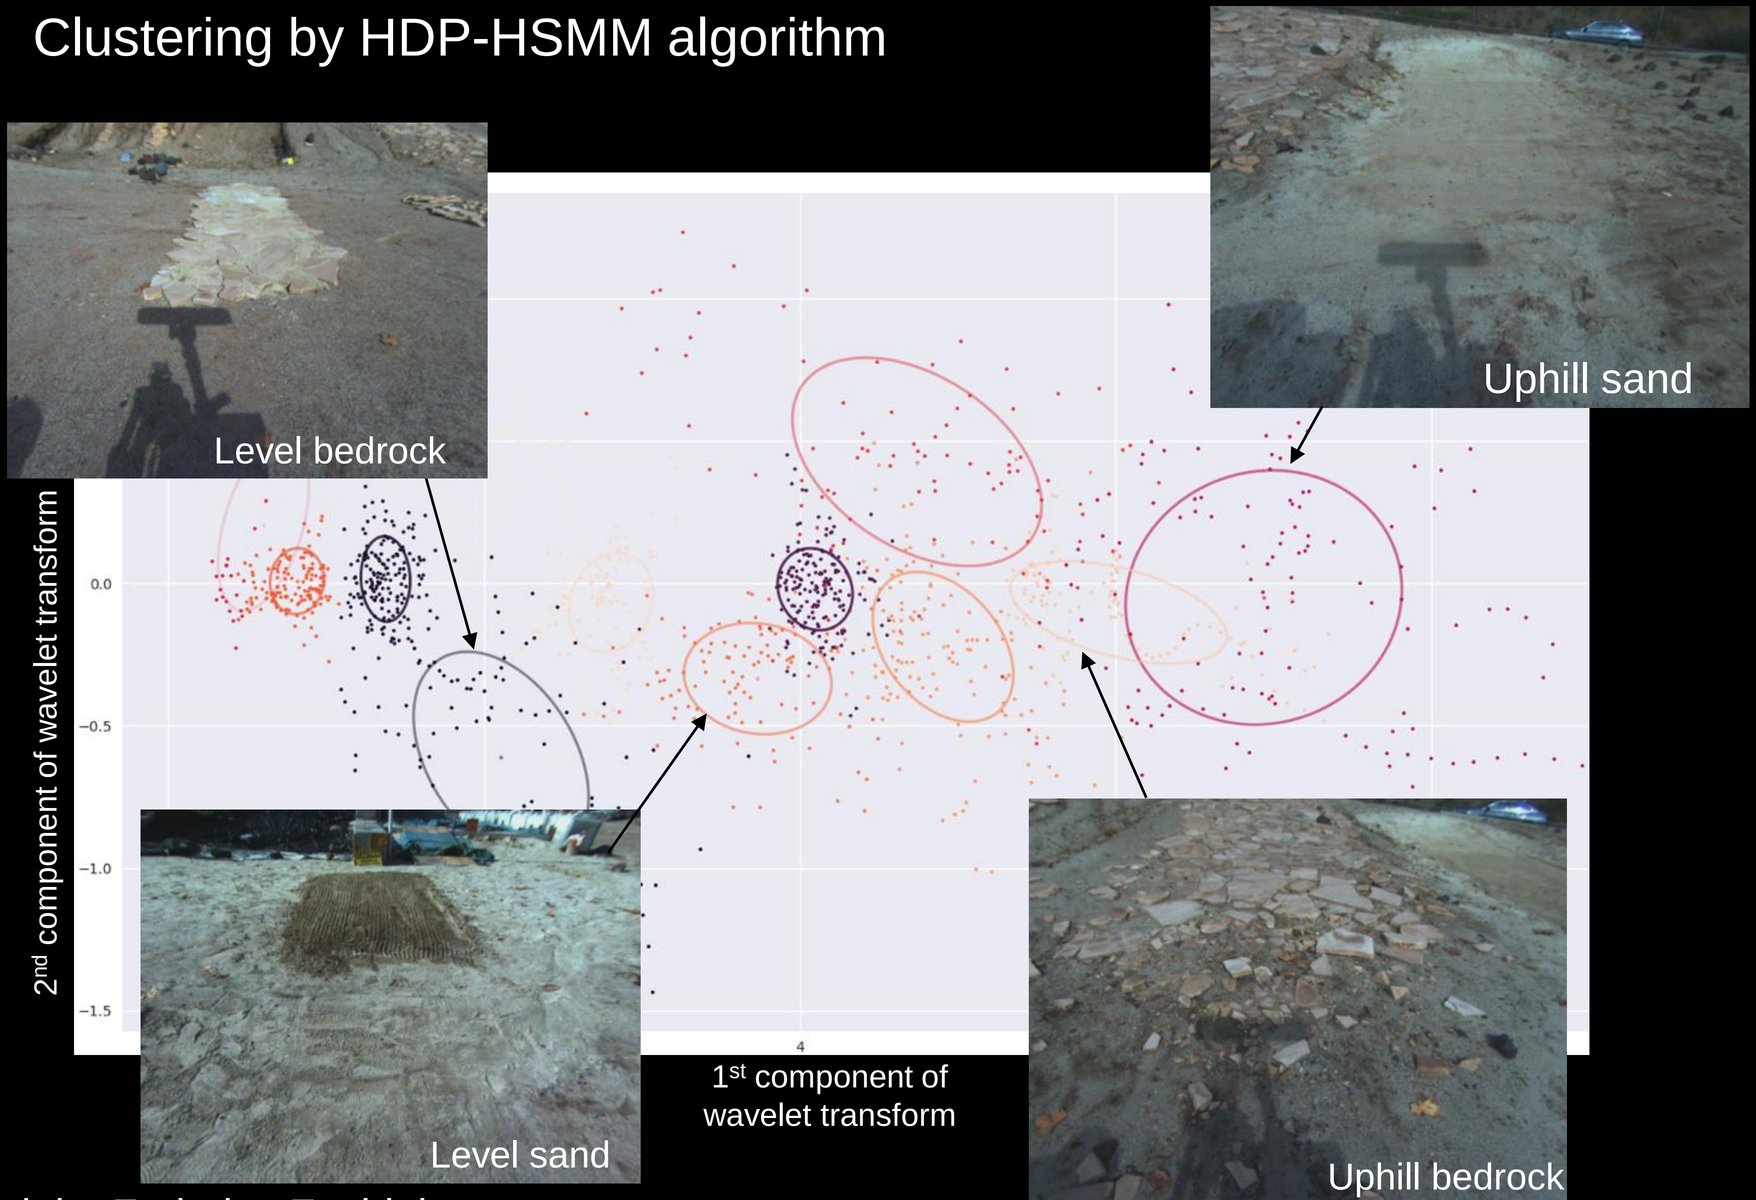

Work by Federico Zechini

Single Wheel Testbed

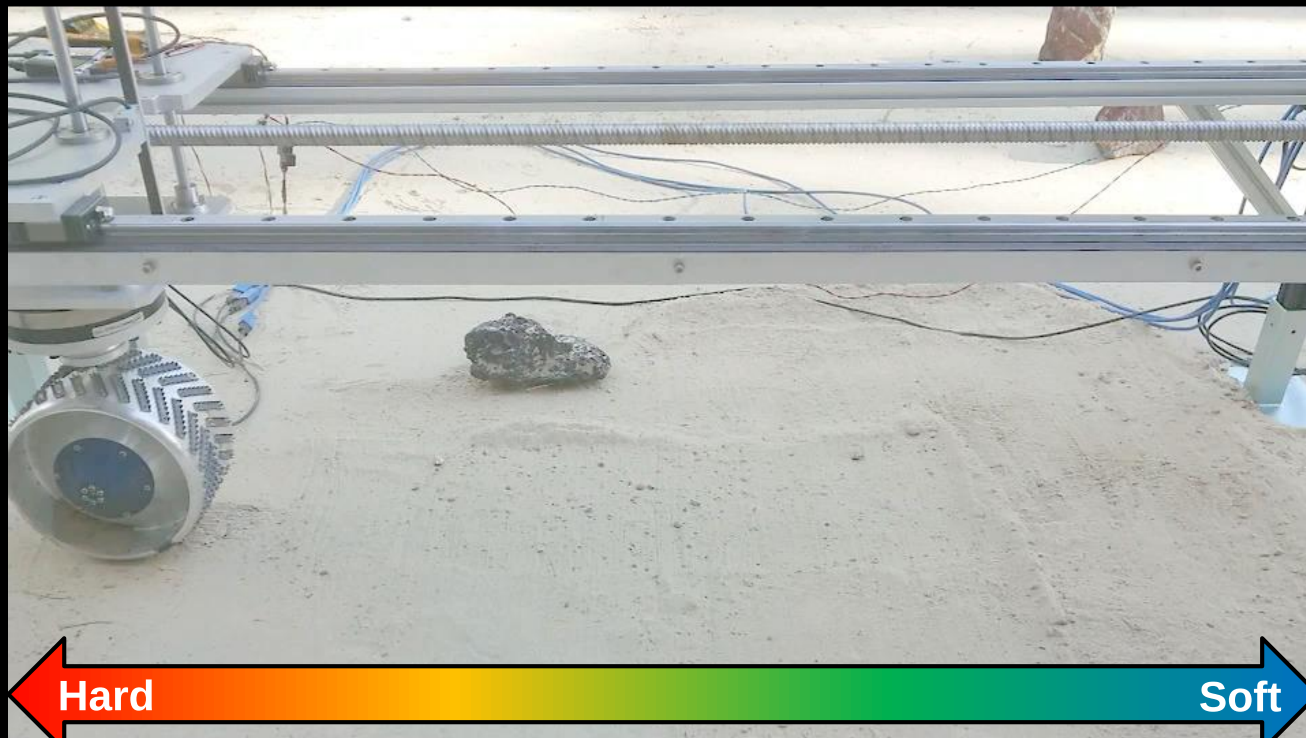

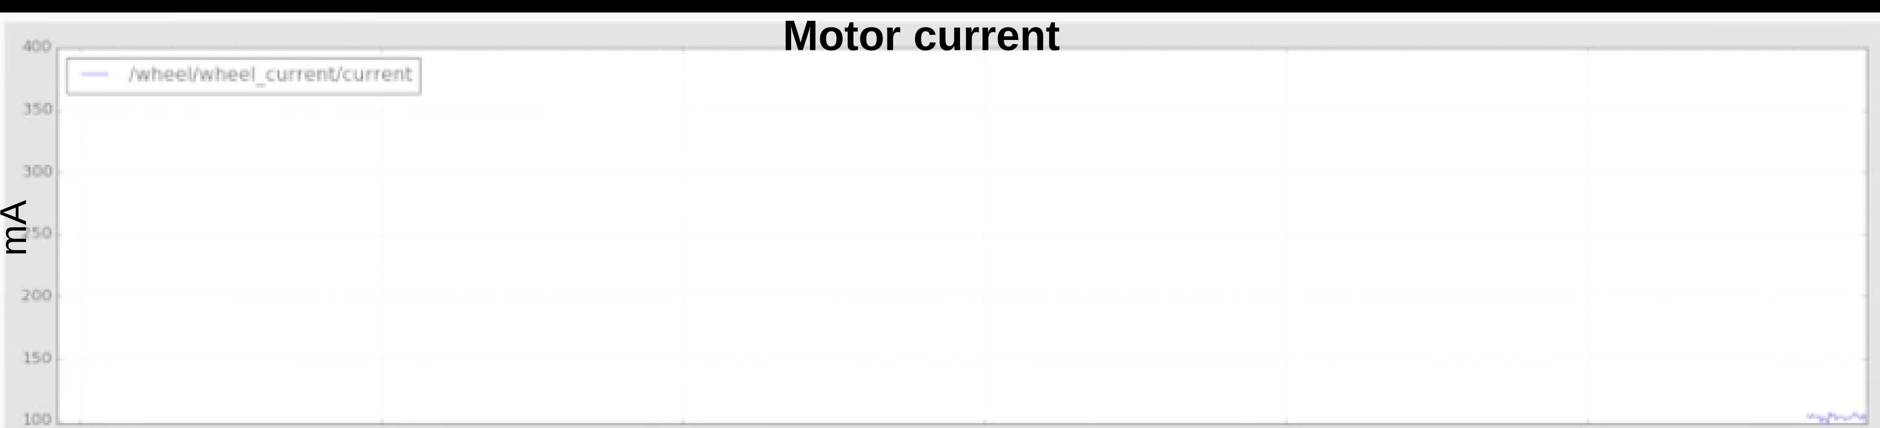

Work by Shoya Higa, Jacek Sawoniewicz

IR-based Terrain Classification: Proof-of-concept

- Created two types of sandy area in Mars Yard:
  - Compact (~80 kPa) and Soft (~30 kPa)
- Measured temperature and soil pressure at 30 locations
- Temperature was collected from 6:30am to 7 pm

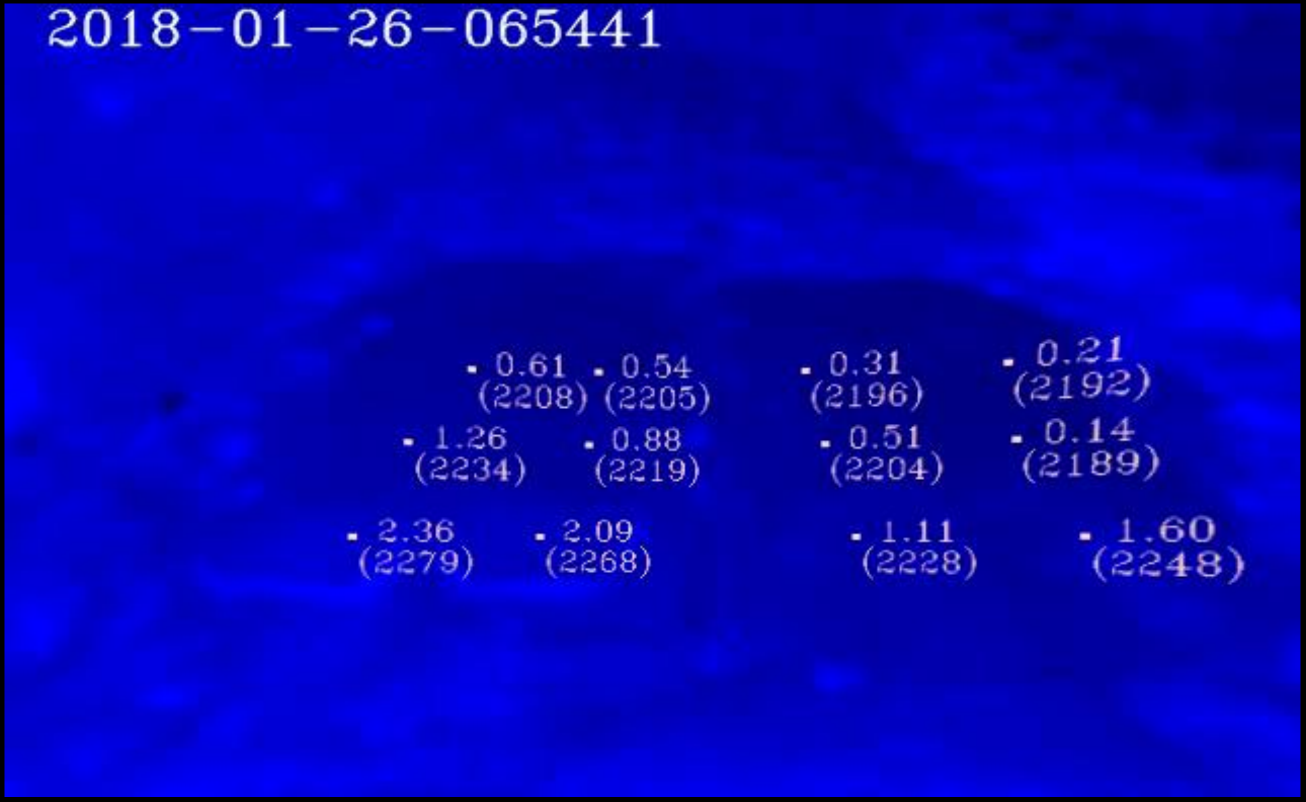

Work by Yumi Iwashita

Demonstrating HPSC

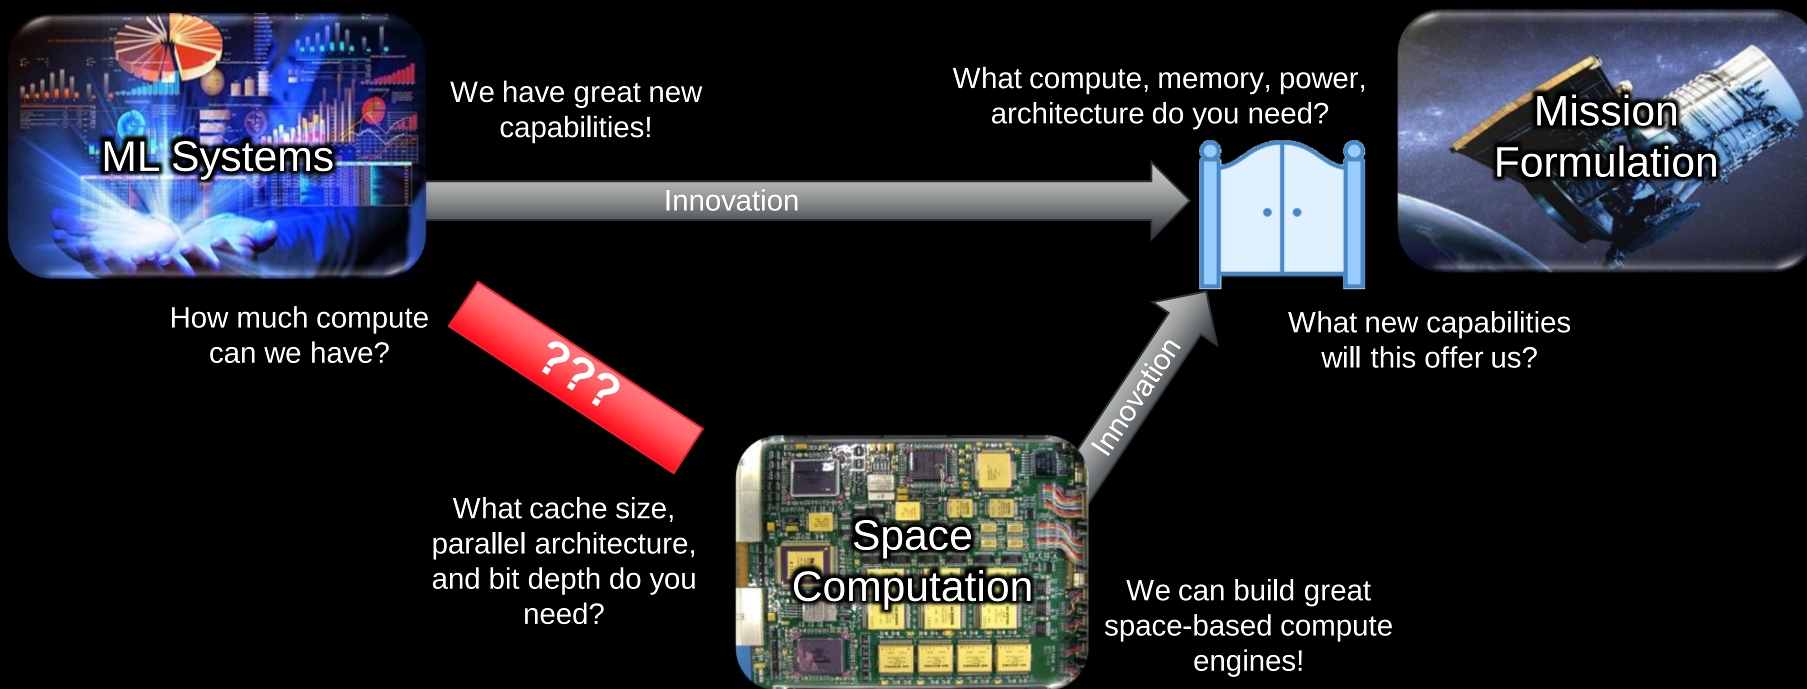

Bridging The Gap: Actual Performance Metrics

Demonstrating HPSC

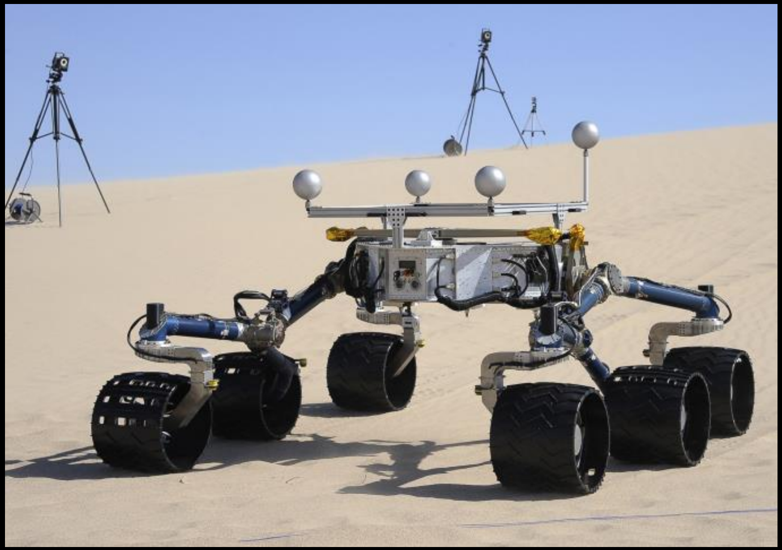

Unique Engineering Sensors

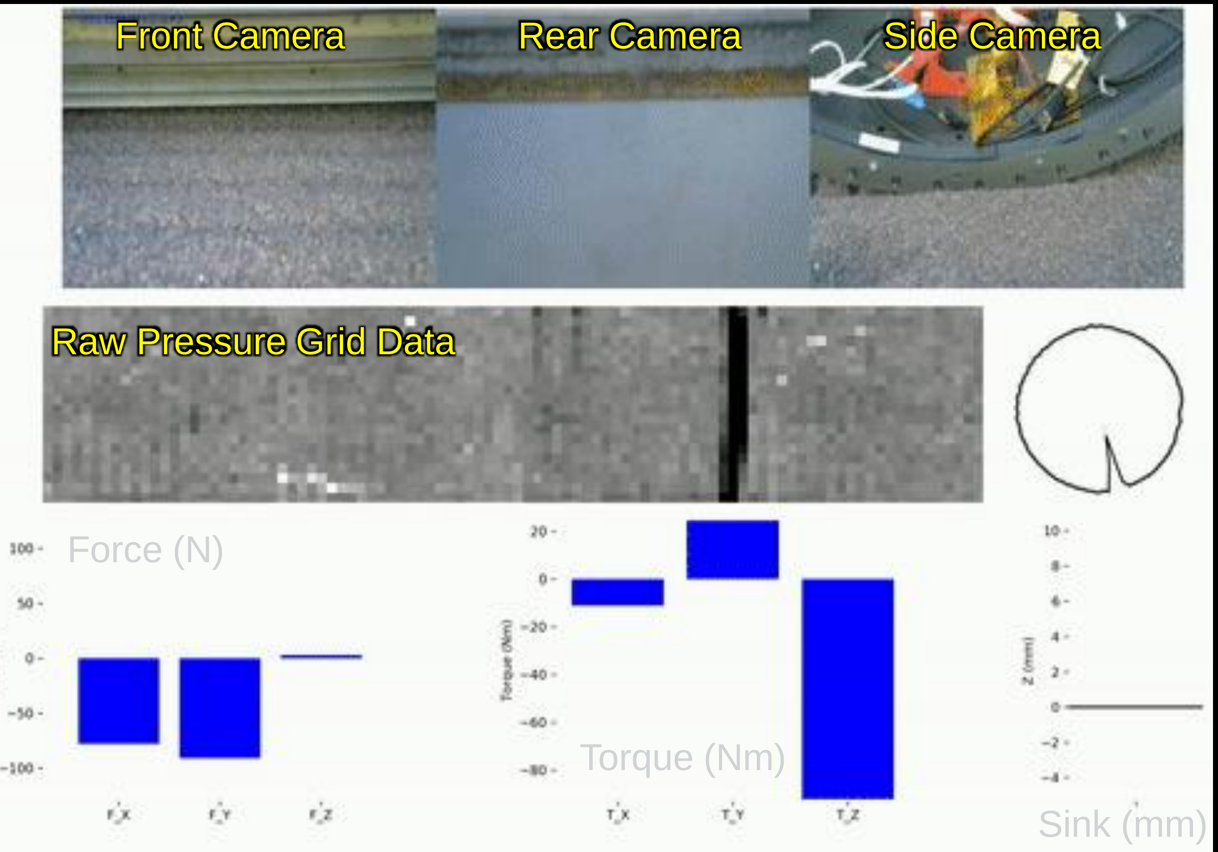

Sensors reporting

- Context cameras
- Pressure Grid
- Force/Torque
- Vertical Displacement
- Optical Flow
- IMU (accelerations)

Jet Propulsion Laboratory
California Institute of Technology

jpl.nasa.gov